MAUI HIGH PERFORMANCE COMPUTING CENTER

APPLICATION BRIEFS 2005

A Center of the Air Force Research Laboratory
Managed by the University of Hawaii

MAUI HIGH PERFORMANCE COMPUTING CENTER

APPLICATION BRIEFS 2005

Air Force Maui Optical & Supercomputing Site (AMOS)

MAUI HIGH PERFORMANCE COMPUTING CENTER

550 Lipoa Parkway, Kihei-Maui, HI 96753
(808) 879-5077 • Fax: (808) 879-5018
E-mail: info@mhpcc.hpc.mil
URL: www.mhpcc.hpc.mil

An Air Force Research Laboratory Center Managed by the University of Hawaii.

WELCOME

This is the eleventh annual edition of Maui High Performance Computing Center's (MHPCC) *Application Briefs* which highlights some of the successes our staff and customers have achieved this year.

MHPCC, established in September 1993, is an Air Force Research Laboratory (AFRL) Center managed by the University of Hawaii. A leader in scalable parallel computing technologies, MHPCC is primarily chartered to support the Department of Defense (DoD) and other government organizations.

MHPCC offers an innovative environment for High Performance Computing (HPC) applications. This includes:

- Computational Resources: Stable and secure parallel computing platforms for prototyping, benchmarking, and testing applications. MHPCC is ranked as one of the top HPC centers in the Department of Defense in terms of computational capabilities.

- High-Speed Communications Infrastructure: OC12 connections, offering 620 megabit per second (Mbps) capacity, provide direct access to MHPCC resources — over the Defense Research and Engineering Network (DREN) and the Hawaii Intranet Consortium (HIC).

- Support Services: An expert staff provides MHPCC users with systems, network, and applications support in addition to assistance with code porting, optimization, and application development.

MHPCC is a well-established member of the High Performance Computing community, participating in collaborations and partnerships that extend its capabilities. MHPCC is a direct contributor to the Department of Defense as a:

- Center within the Air Force Research Laboratory. MHPCC works closely with DoD and other government researchers to support Research, Development, Testing, and Evaluation (RDT&E) efforts.

- Allocated Distributed Center within the DoD High Performance Computing Modernization Program (HPCMP). MHPCC provides resources to the DoD research community, as well as Pacific Region DoD organizations, including the Air Force's Maui Space Surveillance Complex.

- Air Force Research Laboratory resource for the Air Force Maui Optical & Supercomputing Site (AMOS).

- Member of Hawaii's growing science and technology community.

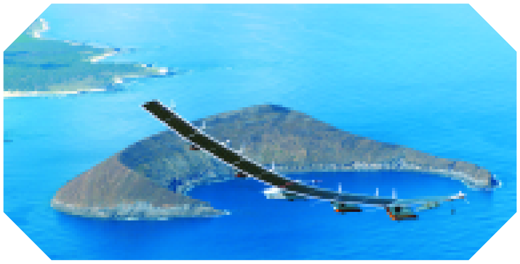

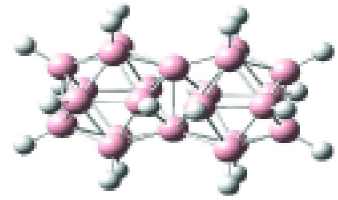

APPLICATION BRIEFS


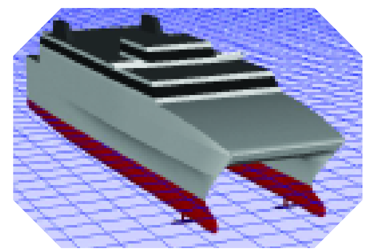

The user application briefs in this publication represent selected research efforts that have taken place at MHPCC during 2005. Each application brief was written by an individual researcher or research team, and reflects their first-hand experiences using MHPCC resources. These articles reflect the diverse nature of our users and projects.

The Application Briefs in this document are the result of the efforts of more than 50 authors. We acknowledge the contributions of each of these individuals and are grateful for their work. We welcome back those authors who have become regular and frequent contributors. We also welcome those making their MHPCC Application Briefs debut this year.

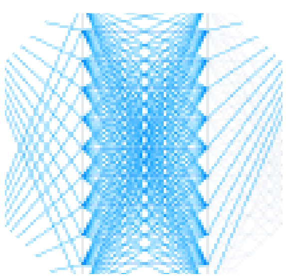
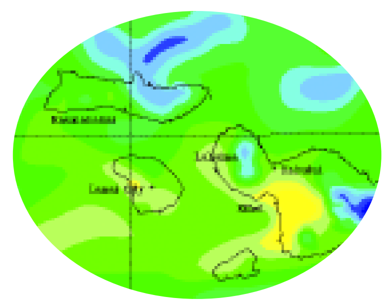

The shaded box at the top of each brief's first page is a short summary of the article. Author and/or organizational contact information can be found in the shaded box at the end of each brief. The notation at the bottom of each page indicates each project's primary functional area (DoD, Government, or Academic).

And finally, feedback regarding this publication is solicited. Please direct any communications to: *editor@mhpcc.hpc.mil.*

Thank you for your support.

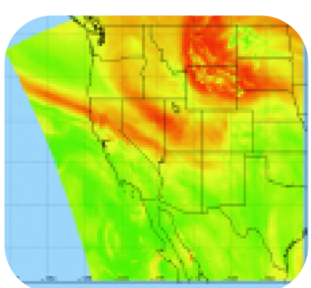

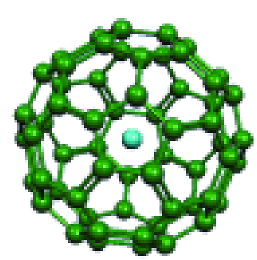

TABLE OF CONTENTS

TABLE OF CONTENTS CONTINUED

Another Level of Parallelism for the Physically Constrained Iterative Deconvolution (PCID) Algorithm

Kathy Schulze and Charles Matson

The imaging mission of the Air Force Maui Optical & Supercomputing Site (AMOS) requires high resolution imagery. Linear image processing algorithms, like bispectrum, used to reduce the observatory data are fast, easily parallelized, and produce a recovery in seconds. Bispectrum cannot recover the very high spatial frequency data we seek due to the direct calculations making up the heart of the algorithm. The high performance computing assets at Maui High Performance Computing Center (MHPCC) make possible the integration of non-linear iterative algorithms as data reduction tools which can recover the very high spatial frequency information. The physically constrained iterative deconvolution (PCID) algorithm, selected as the next generation image processing tool for observatory data, can take minutes to hours to complete a recovery. Our goal is to utilize the high performance computing assets at MHPCC to reduce the execution time of PCID recoveries.

Background: PCID uses ensembles of short exposure imagery containing one or more frames. Each frame corrupted with turbulence from the atmospheric layer of the earth (Figure 1). The algorithm iteratively solves for the common object in each frame, thereby removing the atmospheric corruption. During the processing, the cost function and the gradient function subroutines within PCID are executed many times making heavy use of fast Fourier transforms (FFT).

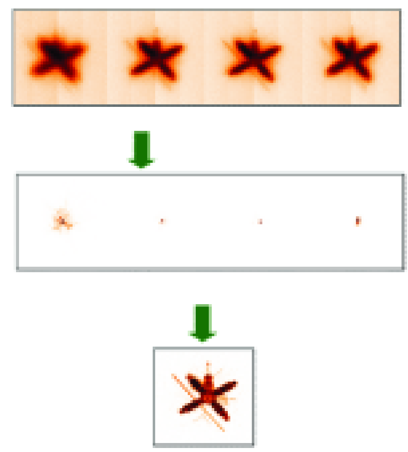

Figure 1. Top row, raw data frames, middle row, PSF bottom row, recovered object.

Our Approach: PCID is capable of handling one or many frames to form a recovery. Initial versions of parallel PCID spread the frames in the ensemble across multiple processors working together. As shown in the following flowchart, (Figure 2) worker processors under control of the master processor are given entire frames within the ensemble. In this scenario, no more workers than number of frames in the ensemble can be used to reduce the data. With the advent of newer, faster, and more processors available for data reduction, we want to split the data frames across multiple workers to decrease execution time.

We want our new PCID paradigm to operate on the architectures of today and scale well for the architectures of tomorrow. This means PCID must work on shared memory processors (SMP) as well as distributed memory processors (DMP) and processors working together on a recovery may be a combination of SMP and DMP.

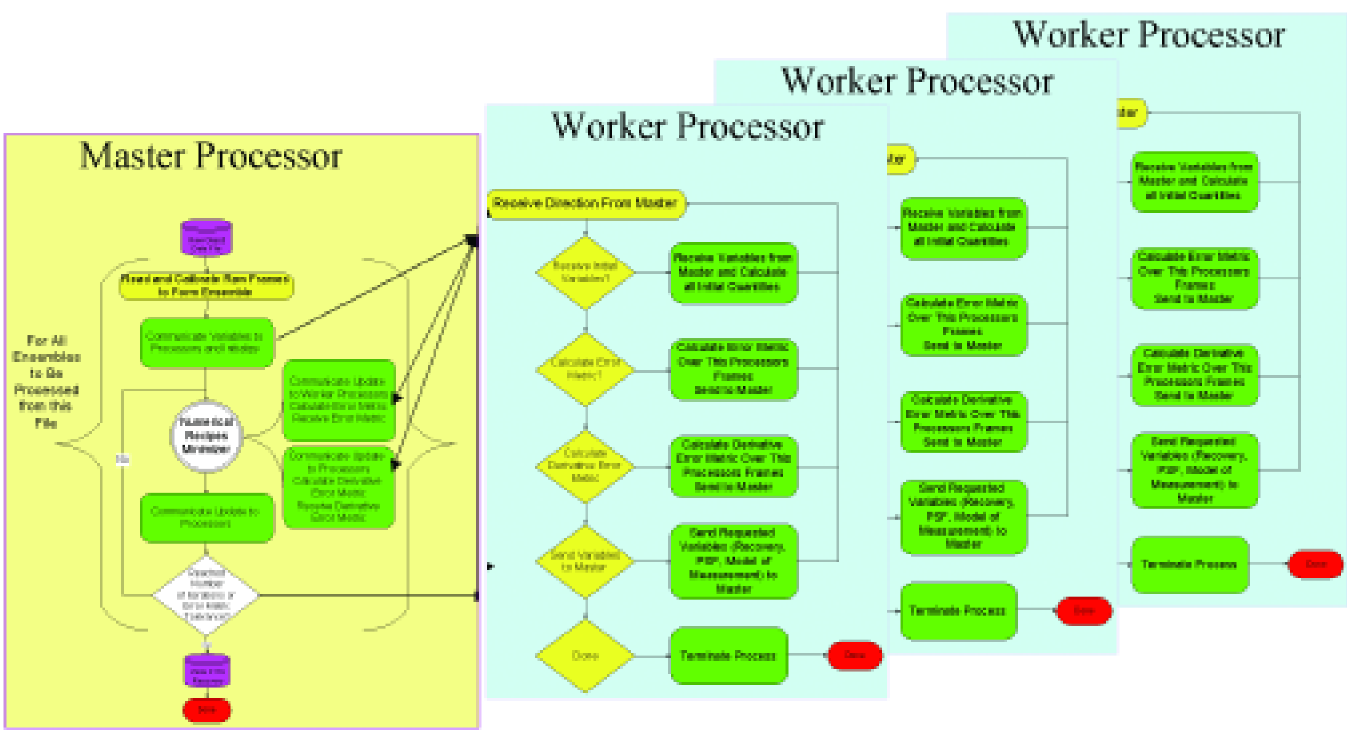

Figure 2. Parallel SMP PCID Flowchart - Whole frame resides on one worker processor.

The cost function and gradient routines require convolutions that are pixel by pixel operations in the Fourier domain. Pixel by pixel operations lend themselves ideally to splitting a frame of data across multiple workers requiring no inter-process communication. FFT operations do require inter-process communication but library functions exist to take care of this process for us. On the IBM *Tempest* system, the Parallel Engineering and Scientific Subroutine Library (pESSL) math library contains FFT routines to accommodate data across multiple processors. Basic Linear Algebra Communication Subroutines (BLACS) contains routines to setup the worker processors so they can be used together for the FFT. BLACS and pESSL work transparently with SMP and DMP processors. BLACS is portable and, in general, each new system has its own parallel optimized math library comparable to pESSL.

The first step was to understand and control the parallel FFT routines. A stand-alone program was written (outside of PCID) to test the FFT and get some timing stats, which we hoped would indicate the level of performance improvement we could gain with single frame, multi-processor PCID. Our initial results from the stand-alone FFT test routine indicated a speedup that correlates with the number of added worker processors.

The second step was to modify PCID to utilize multiple worker processors for a single frame of data. Processors working together on a single frame of data are known as a processing set and are shown in the flowchart here (Figure 3).

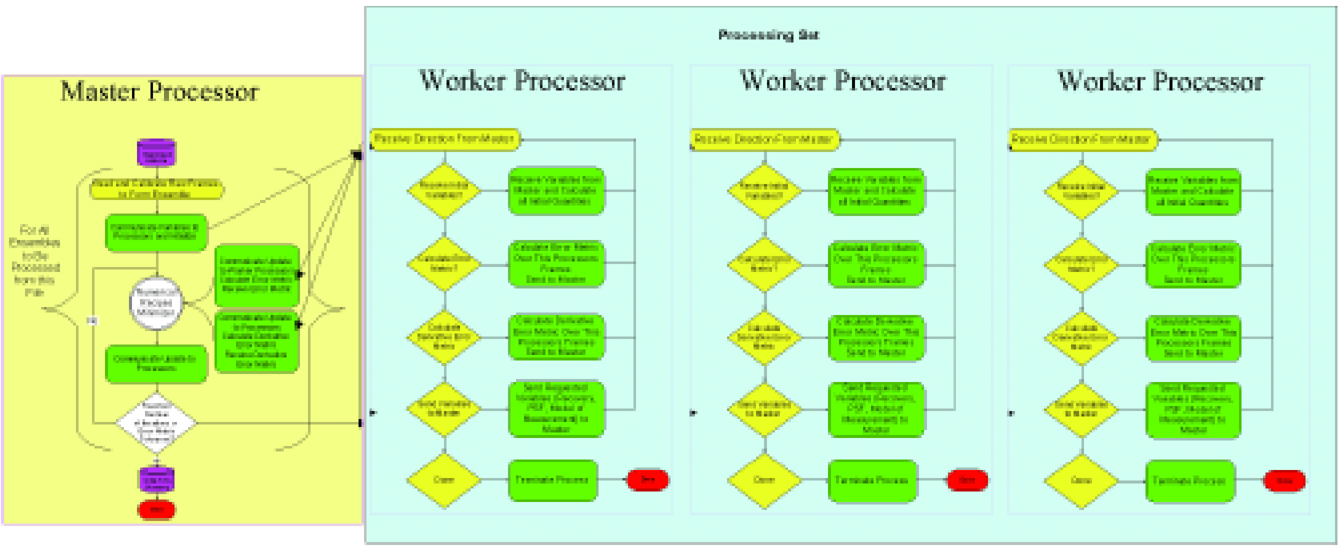

Figure 3. Parallel DMP PCID Flowchart - Frame split across worker processors.

- Each worker handles a portion of a frame.
- Each process set handles a complete frame.

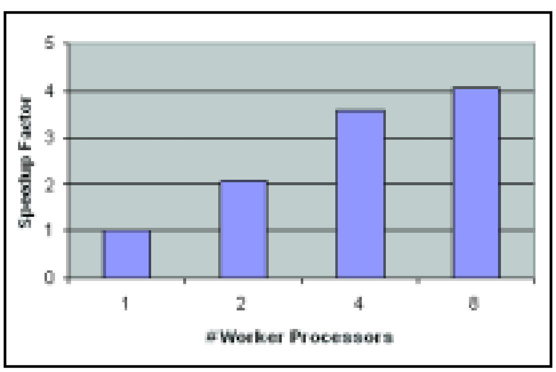

Figure 4. Multi-processor Single Frame PCID.

We see in this plot that in the multi-worker, single frame case, we are achieving our expected speedup correlating with the number of added worker processors. (Figure 4).

Summary and Future Work: We show that splitting a frame of data across multiple worker processors does decrease execution time and the decrease correlates nicely with the parallel FFT results. This correlation is completely expected since the cost and gradient functions in PCID are heavily dependent upon Fourier transforms.

Our future work includes processing sets containing multiple frames and multiple processing sets, to reduce a many frame ensemble of imagery to produce one recovery.

Author and Contact: Kathy Schulze
Organization: KJS Consulting, 71 Awalau Road, Haiku, HI, 96708
Author: Charles Matson
Organization: Air Force Research Laboratory/DESA, Kirtland AFB, NM, 87117-5776

Resources: *Tempest* P3, P4 at MHPCC

Sponsorship: Air Force Research Laboratory

High Performance Computing Software Applications Institute for Space Situation Awareness (HSAI SSA)

Bruce Duncan, Kathy Schulze, Mike Berning, Kevin Roe, Aaron Culliney, Robert DeSonia, Eric Van Alst

The DOD High Performance Computing Modernization Program (HPCMP) has initiated a High Performance Computing Software Applications Institute (HSAI) project which was started in the government fiscal year 2005. The HPCMP philosophy is to evolve its original Computational Technology Areas (CTA)-based program, with emphasis on software testing, to an integrated computational science and engineering applications development and support program that focuses on the sum of all previous technologies to address chosen major DOD needs and programs.

Background: The DOD High Performance Computing Modernization Office (HPCMO) originally devised the strategy for the HSAI in 2003 and issued a call for DOD organizations proposals late that year. After initial vetting stages by each of the military services, 12 organizations were ultimately invited to submit proposals to the HPCMO. HPCMO made and announced its final selection of five organizations on 20 August 2004 to stand-up Software Application Institutes (SAIs), as shown in Figure 1. One of the five institutes that were selected by the Deputy Under Secretary of Defense (Science and Technology) for activation was an HSAI for Space Situational Awareness (HSAI SSA), which was awarded to the Air Force Research Laboratory's Directed Energy directorate (AFRL/DE). The HPCMP Kick-Off Meeting for the five institutes occurred in October and initial partial funding was made available in December 2004. The first specific meeting of the HSAI SSA with the HPCMO's Board of Directors occurred in February 2005. At that time, the HSAI SSA briefed its approach to structuring and staffing of the HSAI SSA and strategic goals and technology thrust areas, with particular emphasis on activities to be performed during the first year of the program. The HSAI SSA continued further program initiation work during March 2005.

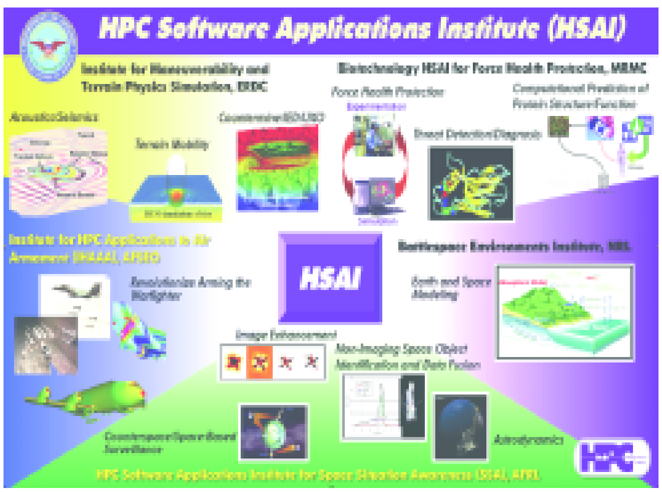

Figure 1. HPC Software Applications Institutes.

Methodology: The mission of the HSAI SSA is to support the Space Situational Awareness needs of stakeholders by developing high performance computing (HPC) software applications for SSA. SSA encompasses the Space Support and Mission Support foundation tiers of United States military space power. The Institute will apply the power of high performance computing and innovative algorithms to enable the warfighter to better perform space control missions by providing a more accurate, more detailed, and more timely characterization of space objects, and by developing a data management architecture that can be utilized by the stakeholders. Early HSAI work incorporated program planning and control activities including the establishment of the organization and partnering and supporting organizations, and initial programmatic and technical liaison.

The Air Force established a Task Order (TO) on the existing AFRL MHPCC Contract to support all necessary Institute initial requirements. This avoided the need for establishing a separate contract and enabled more rapid securing of personnel resources for early software management and development work, and it arranged for the securing of office space and provision of communications and other infrastructure.

HSAI SSA has refined its original planning and has chosen two Strategic Goals and technical thrust areas to begin its work in. These items are Strategic Goal 1: Astrodynamics for the Characterization of Space Objects, and Strategic Goal 2: Image Enhancement.

The major objective of the Astrodynamics thrust is to develop advanced orbit prediction and metric extraction capabilities. This involves:

- Development of HPC Application Software to Aid Space Surveillance Network (SSN) Architecture, Systems and Sensors Developments, and Operational Decision-Making Processes
- Expansion of Spacetrack Test Bed Architecture for Air Force Space Surveillance System (AFSSS) Fence Radar Developments, Architecture, and Operational Decision-Making Processes
- Development of a Parallelized Scalable Satellite Orbit Propagation Tool (Parallel Catalog Propagation - PCP) for Higher Accuracy Orbit Calculations

The major objective of the Image Enhancement area is to derive enhanced utility from electro-optical sensors by applying various image processing algorithms currently in development for SSA solutions and enhancing them further to better meet customer requirements. Efforts here are focused toward investigating and resolving the issues with the implementation of the current and new algorithms in the HPC environment. The HSAI SSA team has begun the process of developing a prototype that will form the foundation for an operational product. Customer and stakeholder organizations for these items include the U.S. Strategic Command (STRATCOM), Air Force Space Command (AFSPC), and Air Force Maui Optical & Supercomputing Site (AMOS). This involves:

- Expansion and Implementation of New Advanced Image Processing Software and Data Management Tools in the HPC Environment
- Enhancement of the Fidelity of Image Products and Development of Scalable High-Speed Parallel Versions Using HPC Assets to Decrease Computational Times
- Employment of the Physically Constrained Iterative Deconvolution (PCID) Software and Multiple Instruction Multiple Data (MIMD) and Multiple Process Multiple Data (MPMD) Computational Methods to Efficiently Yield Consistent, High Resolution Images

An example of early work performed by HSAI SSA includes recent efforts aimed at providing higher resolution, more precise space object imagery to the U.S. military. HSAI SSA modified previously developed PCID application source code so that multiple processors can be used to carry out cost function and cost function gradient calculations for a single frame of data. This software is comprised of more than 16,000 SLOC. The team demonstrated more than a two-fold improvement in processing speed (see Figure 2) over the current version of PCID through the software engineering and code optimization that was performed. A validated, demonstrated, and fully successful metric was achieved for this goal. Furthermore, this goal was achieved for multi-frames, multi-processors, and multi-processing sets for distributed memory and shared memory HPC systems utilizing MHPCC *Tempest* IBM P3 and P4 processors.

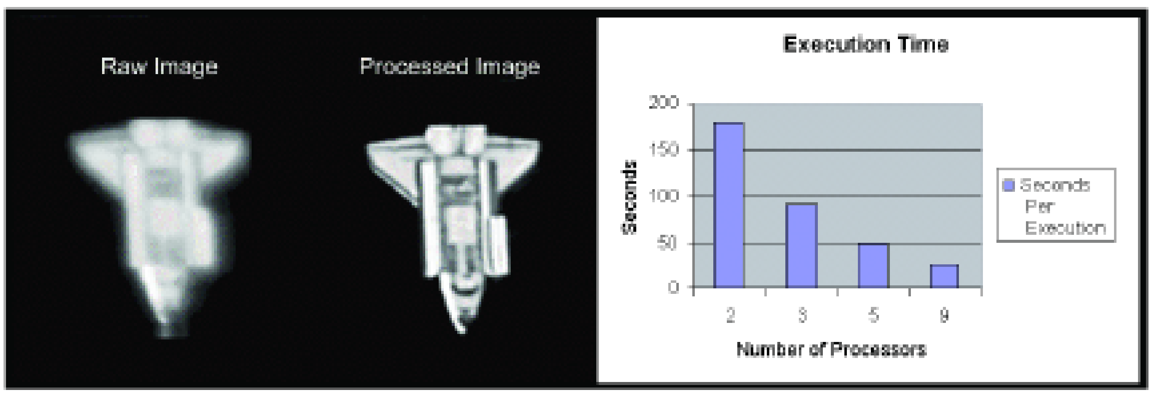

Figure 2. HPC PCID Provides Higher Resolution Space Object Images for DoD, with Significant Speedup.

Similarly, additional work was also performed on SSNAM software to evaluate scalability, timing, sizing, and differences in running this code on different platforms. Initial work was performed completing porting and integration of more than 560,000 SLOC and setting up for evaluation runs under the IBM AIX P3/P4 MHPCC *Tempest* environment, employing a 64-Bit implementation. The validation and initial evaluation timing runs were as anticipated, demonstrating good scalability, throughput, and showing that a factor of two speedup is attainable. The HSAI SSA team next ported and integrated the executable to the JFCOM J9 *Koa* Linux Networx Evolocity II platform at MHPCC. The team executed the software on a single *Koa* 2.4 GHz Xeon processor. The timing results and evaluation from this activity demonstrated that a factor of eight speedup was achieved against a test case run on a SSNAM technical support contractor's Laptop PC. The results also showed that a factor of three is achievable for the test case run on a single CPU of the SSNAM technical support contractor's SSNAM Laboratory PC (see Figure 3). These results were very encouraging, as follow-on HPC development work will be able to run with significantly more processors than the eighteen that are available in the current SSNAM PC lab. It is estimated that a speedup of five-fold, or greater, is very likely to be realized in the near future.

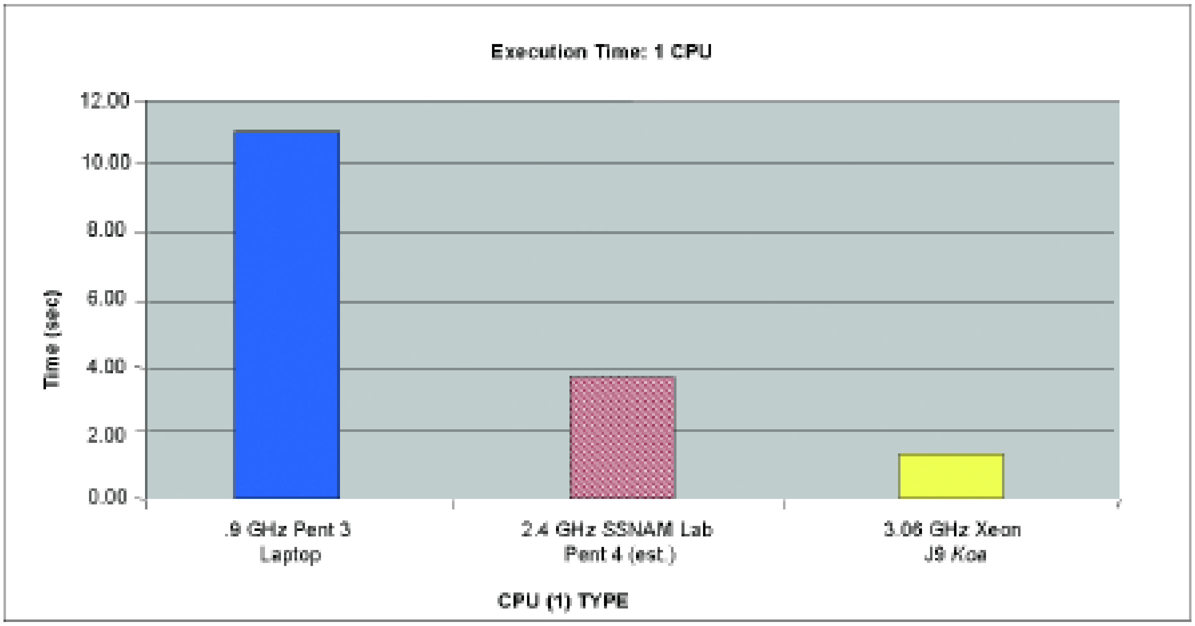

Figure 3. HPC SSNAM Provides Higher Speed Runs to Analyze the Space Surveillance Network of Sensors.

Significance: The HSAI mission and resulting efforts will expand DOD Service/Agency efforts in providing robust production-level software for modeling, simulation, and computation in HPC application areas of the highest impact to DOD. The goals and focus areas of the HSAI SSA have direct applicability to customer requirements and ongoing activities to increase DOD space surveillance and exploitation capabilities that are important to the defense of the nation. The increasing number of satellites (with decreasing structural sizes) and the increased capabilities of space surveillance sensors will necessitate advanced software applications that utilize high performance computing assets. HSAI SSA will exploit technical resources to develop, apply, and support the transition of DOD HPC software applications. This will help lower cost and accelerate electronic systems design, development, testing, and procurement while capitalizing on modeling and simulation where applicable, and provide for a greater usability and robustness of applications software across a wide user base.

Author and Contact: Bruce Duncan
Authors: Kathy Schulze, Mike Berning, Kevin Roe, Aaron Culliney, Robert DeSonia, Eric Van Alst
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Maui, HI, 96753
Resources: IBM P3 and P4 Servers and Linux Networx Evolocity II Supercluster at MHPCC
Sponsorship: Air Force Research Laboratory (AFRL/DESM)

Pan-STARRS Image Processing Pipeline Software Development

Bruce Duncan, Michael Berning, Robert DeSonia, Eric Van Alst

Pan-STARRS (Panoramic Survey Telescope and Rapid Response System) is an innovative new design for a wide-field imaging facility being developed for the Air Force Research Laboratory by the University of Hawaii (UH) Institute for Astronomy (IfA). UH IfA is the prime contractor and leader of the project. Massachusetts Institute of Technology Lincoln Laboratory (MIT/LL) is designing the charge-coupled devices (CCD) for the project. Science Applications International Corporation (SAIC) is a subcontractor to the IfA working on the database aspect of the project, and Maui High Performance Computer Center (MHPCC) is supporting IfA in the development of the data processing pipeline. By using four comparatively small co-located telescopes, the Pan-STARRS team plans to deploy an economical system that will be able to observe the entire available sky several times each month. The near-term goal is to discover and characterize Earth-approaching objects, including asteroids and comets that might pose a danger to our planet. The huge volume of images produced by this system will also provide valuable data for many other kinds of space scientific programs and products. This paper summarizes the MHPCC contractor team development work performed to date.

Research Objectives: The full system, PS4, will be composed of four individual telescopes of 1.8 meter aperture observing the same region of sky simultaneously. Each telescope will have a three-degree field of view and be equipped with a CCD focal plane mosaic with more than one billion pixels. Pan-STARRS will cover 6,000 degrees squared per night in the survey mode and search for potential killer asteroids. The whole available sky as seen from Hawaii will be observed three times during the dark time in each lunation. Camera exposure times will vary between 30 and 60 seconds and IfA anticipates that Pan-STARRS will reach a limiting magnitude of 24. The focal plane will employ orthogonal transfer charge coupled devices (OTCCDs) that allow the shifting of charge along rows and columns, thus providing on-chip image motion compensation that is the equivalent of traditional "tip-tilt" image compensation, but without moving parts. Each raw image from a single Pan-STARRS camera will contain two Gigabytes of data and it is estimated that the raw data rate will be several Terabytes per night for the full telescope. IfA decided in December 2003 to develop a one-telescope prototype system, PS1, which will be essentially one quarter of the total system and will be completed ahead of the full PS4 observatory. PS1 will have the same optics design and camera design as anticipated for the full version of Pan-STARRS. PS1 is being constructed inside a building that previously housed the Lunar Ranging Experiment (LURE) observatory atop Haleakala, Maui. A new dome will be placed on the building. Figure 1 shows an aerial view of the PS1 site, while Figure 2 shows the PS1 telescope in its dome. PS1 will allow the team to test all of the technology and subsystems that are being developed, including the telescope design, the cameras, and the data processing and reduction software. PS1 will be used to make a full-sky survey that will provide astrometric and photometric calibration data that will be used for the full Pan-STARRS survey. First light for PS1 is scheduled for early calendar year 2006, with deployment of the full PS4 system in approximately two additional years. Sites on Mauna Kea and Haleakala are being considered for the PS4 system, and site evaluations have been performed in parallel at both locations.

Figure 1. PS1 Installation in former LURE facility. PS1 telescope and dome will replace existing dome on left side.

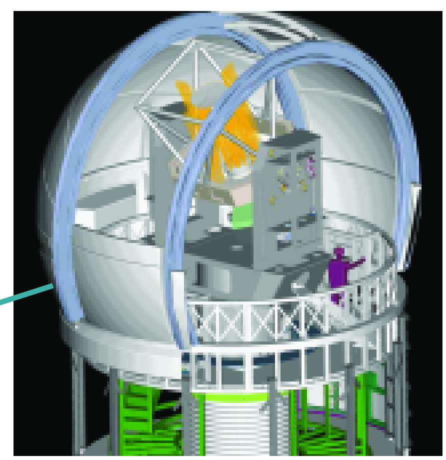

Figure 2. PS1 Telescope Dome installation pictorial (UH IfA).

Methodology: The Image Processing Pipeline (IPP) team was committed to executing the most conscientious software development program as possible at the outset of the PS1 build phase. This phase began in April 2004. The primary objective was to develop a sound product, utilizing software development techniques and tools which provide the best return on investment (labor and capital expenditures) without compromising quality. The key parts of the PS1 IPP development cycle plan include program planning and control, software planning and management, software requirements and architecture, software detailed design, coding, and unit testing. Future team plans beyond the period covered by this writing include unit integration testing, qualification testing, hardware and software integration testing, system qualification testing, and operations. The Pan-STARRS IPP software development team was organized such that UH IfA personnel had the overall lead for the work, and they developed the software requirements and algorithms definition for each build cycle. The MHPCC contractor personnel wrote Software Design Description documentation, reviewed and helped refine algorithms as necessary, and performed coding, unit testing, and quality assurance of the associated source code and documentation. Each software delivery made by the MHPCC team included a Software Version Description document, which accurately described the requirements satisfied by the delivery and the results of internal testing. Doxygen and ManPage programming documentation were also provided. Other standards, methods, and tools employed on the project include:

- IEEE 12207 and MIL-STD-498 Systems and Software Engineering Practices and Disciplined Team Software Process Regimen
- ANSI 'C' Programming Language and Callable Functions Implementation
- Perl for Flow Control and High Level Image Processing Functions
- Object-Oriented Design (OOD) model accomplished in C via Coding conventions
    - Object's data attributes are grouped into a C struct arrangement and the associated methods are named with the struct name as a prefix
    - Software hierarchy:
        - psScripts
        - psModules
        - psLib (Library)
        - Industry standard applications/utilities, such as GNU Compiler Collection (GCC), 'C' Flexible Image Transport System Input Output (CFITSIO), Fastest Fourier Transform in the West (FFTW) library, and the GNU Scientific Library (GSL)
        - Glib C Library and GNU C extension (compiled by gcc)
- Linux x86 and Mac OSX compatibility
- Portable Operating System Interface (POSIX)
- Compliance with Industry and Pan-STARRS Project Coding Standards
- Concurrent Version System (CVS) - Software version control
- Bugzilla - Defect and Bug tracking and disposition
- GForge - Software collaboration/management/tracking/test point generation tool
- Software metrics tracking, including Defect/Bug classification and status, Requirements Verification, and Test Points Tracking, etc.

Results: Deliveries of the first six releases of PS1 IPP software represented on-time deliveries and substantial milestones in the progression of the overall Pan-STARRS project, while increasing our team software process acumen. The most recent release prior to the publication timing requirements for this paper was Release 6.0, which was delivered on 13 June 2005. This was the sixth of at least twelve builds that have been planned by IfA. A high level profile of the metrics for Release 6.0 is outlined below.

- 72, 295 Source Lines of Code (SLOC, with comments)
- 704 Total psLib and Modules Requirements Verified
- 188 psLib and Modules Test Points Verified
- 146 Pages of Technical Documentation: Software Design Description and Software Version Descriptions; and additional Doxygen documentation

Significance: The Pan-STARRS system will provide unprecedented capabilities to the astronomical and scientific communities when completed. The MHPCC contractor team is continuing to support IfA by developing application software source code and documentation that is at the very core of the Pan-STARRS system. The project team will continue working diligently toward the PS1 prototype system, with the longer-term goal of supporting the fielding of the full four-telescope PS4 system.

Author and Contact: Bruce Duncan
Authors: Michael Berning, Robert DeSonia, Eric Van Alst
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Maui, HI, 96753
Resources: IBM Linux Supercluster, and Software Development Servers at MHPCC
Sponsorship: Air Force Research Laboratory
Acknowledgement: The MHPCC contractor team wishes to express its thanks to UH IfA for its pioneering work and leadership on Pan-STARRS, and its support of the MHPCC team.

Millimeter Wavelength Imaging of Orbiting Satellites

Bobby Hunt and Robert R. Borchers

From the success of radio astronomers in producing images of the remotest objects in the universe it is logical to raise the question: is it possible to utilize the same technology to make images of satellites in earth orbit? Answering this question is the purpose of a Directed Technical Task at the Maui High Performance Computing Center (MHPCC) during the past year.

Introduction: To monitor orbiting satellites the Air Force maintains state-of-the-art telescope facilities at the top of Mt. Haleakala on Maui. These telescopes make images of objects orbiting the earth, and detect and track satellites passing overhead. Collecting images by light has drawbacks that are familiar to every person, however. Clouds and other weather conditions often block viewing of the sky. Even in the Hawaiian islands, with some of the best weather conditions found anywhere on Earth, a typical year will see the view of the sky, by the telescopes on Mt. Haleakala, blocked as much as one-third of the time. Although light waves are not able to pass through clouds, radio waves travel through clouds with little effect. The ability of radio waves to convey information in the presence of clouds is also seen in the practice of the most computational intensive method of astronomy: radio astronomy. Radio astronomy makes images of the universe by using large antennas to capture radio waves emitted by interstellar, intergalalctic, and extragalactic objects. The recordings of the radio signals are then converted to images by special computer complexes. Adapting the technology of radio astronomy to the imaging of earth satellites would produce a capability for collecting satellite information that would not be seriously affected by weather.

Methodology: One of the important contributions of Einstein was the photon. A consequence of the existence of photons was the discovery that energy is radiated from all objects that are any warmer than absolute zero. The Boltzmann radiation equation governs this emission of energy, and is often referred to as the "black-body" radiation law. The Boltzmann radiation equation predicts radiation at electromagnetic wavelengths normally associated with radio. This type of radiation is referred to as "passive," because there is no requirement for an active source of controlled radiation.

In order to make pictures of the objects that emitted the radio waves collected in the antennas of a radio telescope, it is necessary for the radio signals collected in the antennas to be processed by computer. The computer processing is carried out in a manner that is analogous to the interference of waves that occurs in the mirrors and lenses of an optical telescope, i.e., a telescope that makes images in the wavelengths of electromagnetic energy that are associated with images visible to the human eye. The computer calculations are referred to as interferometric correlation.

The research that is on-going at MHPCC is directed to two of the most essential aspects of the problem of making images from radio waves emitted by satellites: (1) quantifying the exact level of energy available from typical satellites and the relationship of that energy to the noise that occurs in collecting the radio signals; and (2) determining the best computational architecture by which to organize the processing of the signals into an image.

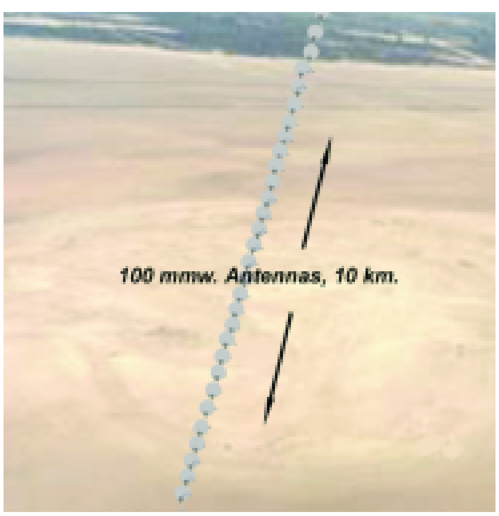

Figure 1. Artist rendition of a linear antenna array.

The investigations of radio signals for forming images of orbiting satellites have focused on signals with wavelengths on the order of a few millimeters (i.e., a few one-thousandths of a meter). This is because of another important property of image formation. The resolution of an image is usually defined in terms of the smallest details that may be perceived in the image. A high-resolution image makes it possible to determine small details and structures of an object seen in the image, and high-resolution images are essential for the purposes to which the U.S. Air Force collects images at the Maui facilities. Basic principles of optics relate the wavelength of energy collected for image formation of an image to the resolution of an image and to the size of the antenna used to collect the image. The resolution of a radio image is directly proportional to the diameter of the antenna and inversely proportional to the wavelength of the collected signals.[1] Thus, to produce a high-resolution image one will seek the largest antenna and collect signals with the smallest wavelength possible. Unfortunately, simple calculations from the optical theorems of diffraction show that using radio signals of the wavelength in ordinary commercial radio, to make high-resolution images of an orbiting satellite, would require the use of an antenna that is thousands of kilometers in diameter.[2]

Reference Notes:
[1]Resolution is conveniently specified in terms of the engineering measure of bandwidth, with the greatest bandwidth being associated with the greatest resolution. Thus, the bandwidth of a radio image increases with an increase in the antenna diameter and a decrease in the wavelength of the collected radio signal.

[2]Of course, we would also not choose to use such wavelengths for making images because of the all the interference from commercial radio that would be present in the images!

However, if we choose radio signals that are of much shorter wavelength, on the order of a few millimeters, then the antenna size can be reduced to a few kilometers. This is also a prohibitively large single antenna, but the use of interferometric image synthesis computations (which we discuss below) makes it possible to achieve this from many small antennas, instead of one large antenna.

When we examine the Boltzmann equation, we discover that the amount of millimeter wavelength radiation from a typical satellite object is quite small. Indeed, the level of millimeter wavelength radiation from a typical satellite is equivalent to that of a cell phone signal from beyond the distance of the planet Pluto! A larger antenna also can mitigate this extremely small level of radiation, since it is a principle of radio engineering that a larger antenna intercepts more energy. However, it is again the case that the size of a single antenna necessary to increase the received energy to a more favorable level is prohibitive.

There are two basic techniques that radio astronomers use to avoid the need to use huge single antennas. The first technique is to use many smaller antennas, placed into an array. The second technique used by radio astronomers is to observe a signal for a very long period. Antenna arrays are not a magic solution, however, because there are costs in the employment of antenna arrays. The signals that are obtained from a single large antenna are directly in a form that can be used to make an image. The signals from an array of antennas are not of this type. The interaction process of waves bouncing off the surface of an antenna is known in electromagnetic optics as interference. It is a fundamental fact of optics that the interference of electromagnetic waves, such as occurring in a single antenna, is responsible for the formation of an image by that antenna. But if two separate antennas are placed at precisely the same places in space as the points on the surface of a large antenna, there is no reflection of the signals to a point of focus beyond the two antennas. Therefore, there is no interference at the focus that is responsible for an image being formed. To make an image, the interference process must be carried out on the signals that are collected from separate antennas.

Understanding the interference process that is responsible for the formation of an image is a matter of writing the equations that describe the propagation of electromagnetic radiation and then introducing into the equations the geometry of an antenna and the way in which signals interact over a period of time. The result of carrying out all the mathematical details of this analysis are found in the Van Cittert-Zernicke Theorem, named after the two physicists who independently discovered the theorem. The salient points of this theorem may be summarized as follows:

- The average, over some interval of time, of the power in two different waves reflected off of two different points of the antenna results in a specific type of computation, known as a correlation. Remarkably, this correlation computation is exactly the type of calculation that a practitioner of signal processing would recognize as the best way to extract weak signals from noise when using long observation periods of the weak signals.

- To synthesize an image from an array of antennas, the signals from all the antennas are collected over long periods of time, and then the correlation computations for the signals from all pairs of antennas are computed.

- Since the process described immediately above is mimicking the way all the signals from all points in a single antenna would simultaneously interfere with each other, the computation is referred to as interferometric correlation for image synthesis.

The calculations to make a radio astronomy image can be carried out by any general-purpose computer, e.g., a desktop personal computer. However, the total number of calculations is so great that no single desktop computer could complete the calculations in any reasonable time. Indeed, even the super computers that are available to the Maui High Performance Computing Center could not keep-up in calculations with the total amount of digital data that a set of radio antennas could collect within in a day.

Radio astronomers have also been faced with the same huge amounts of computation. The response of the radio astronomy community has been the engineering of special-purpose computers referred as digital interferometric correlators. There is currently much innovation in developing new correlators for advanced radio astronomy, the innovation being driven by the emergence of off-the-shelf digital signal processors with great computational capabilities. Using these new digital capabilities the National Radio Astronomy Observatory (NRAO), a facility of the U.S. National Science Foundation that is used by astronomers from all over the world, is currently in the process of designing the next generation of image formation processing correlator. The design of this correlator is still evolving with the development of the new radio astronomy facilities for NRAO, but the current estimates of the correlator performance would make it capable of being adapted for any capabilities needed in satellite imaging by passive millimeter wavelength radiation.

Results: There are a number of engineering and design considerations that need to be resolved for a successful system to make images of orbiting satellites by passive millimeter wavelength radiation, and the study conducted during the last year has been too modest in size to yield any final designs. However, the following summarizes some of the parameters of such a system that our current studies suggest would be feasible and reasonable.

- Radiation wavelength of interest: 94 GHz (approximately 3.2 millimeters) to 240 GHz (1.25 millimeters)
- Diameter of each antenna: 3 meters minimum, up to 12 meters maximum
- Number of antennas: 100 to 1000
- Configuration of antennas: linear array, e.g., all antennas spaced along a single line on the face of the Earth (See Figure 1)
- Antenna spacing: 10 to 100 meters
- Total array dimension: 10 kilometers maximum
- Resolvable object detail for a 10 km Array: Approximately 25 cm maximum (94 GHz), 10 cm minimum (240 GHz)
- Image formation processor: consistent with the coming generation of processors in radio astronomy, i.e.,
    - special-purpose digital correlator, consisting of a many parallel-processing modules;
    - processing modules built on technologies such as off-the-shelf commercial digital signal processing chips developed for advanced video games;
    - total computation capability of several thousand trillion floating point operations per second.

Significance: Besides creating the ability to monitor earth satellites on a 24x7 basis, a successful system using passive millimeter wavelength radiation would offer new methods for studying satellites and developing intelligence on satellites launched by other nations. The different materials used in constructing a satellite have different levels of radiation emission, and study of the emission signatures of satellites in millimeter wavelength images would provide new types of intelligence. Further, the variation in radiation with temperature would provide new operational intelligence in addition to that obtained currently from infrared signatures.

Author and Contact: Bobby Hunt
Author: Robert R. Borchers
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Resources: PCs and Matlab Software at MHPCC
Sponsorship: Air Force Research Laboratory

Terahertz Absorption Properties of X@C_{60} (X = He, Be, F^-, and Br^-)

Tunna Baruah and Mark R. Pederson

Terahertz spectroscopy is fast gaining importance due to its potential applications in communication technology and also in high-resolution imaging and spectroscopy. However, one of the problems facing the development of terahertz lasers is the material source. Practical solid-state sources and detectors do not currently exist for the low terahertz range, also known as the very long wavelength infrared region. One reason is the structural softness of the source materials. Our objective in this work is to explore the applicability of molecular system which show infrared peaks as a source for terahertz lasers. C_{60} is known to possess covalent bonding between the carbon atoms. In this work we present our density functional results on the infrared absorption properties of endohedral X@C_{60} clusters for X = He, Be, F^-, and Br^-. This work may help in identifying terahertz laser sources.

Methodology: The calculations are carried out using the density functional based code Naval Research Laboratory Molecular Orbital Library (NRLMOL) package.[1] The calculations were performed at the all-electron level with the generalized gradient approximation to describe the exchange-correlation effects. NRLMOL employs a Gaussian basis set where the exponentials of the Gaussian basis are optimized for each atom.[2] It also uses a mesh optimized to produce integrals up to a desired accuracy. Further, the vibrational modes were calculated by displacing each atom along the Cartesian axes and calculating the forces for the displaced system. These were used to construct the dynamical matrix and the frequencies were obtained by diagonalizing this matrix.[3]

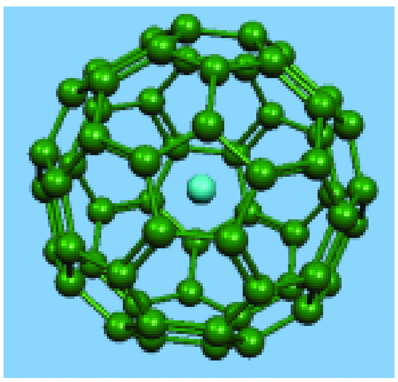

Figure 1. The optimized structures of the endohedral X@C_{60}.

Results: The optimized geometry of the endohedral cluster is shown in Figure 1. The presence of the close shell atom at the center only slightly increases the radius of the C_{60} cage. The HOMO-LUMO gap of C_{60} is 1.64 which remains the same with He but reduces for Be, Br^-, and F^- encapsulated clusters. For the Br and F anions it drops to 0.87 eV and 0.30 eV, respectively. The presence of the closed shell atom is evident in the infrared absorption spectrum. The pure C_{60} shows infrared activity at frequencies 521, 577, 1186, and 1431 cm^{-1}. In the endohedral cluster an IR active modes occur in the frequency range from 100 - 230 cm^{-1}. This mode occurs due to the T_{1u} mode of vibration of the central atom. The IR spectra of the X@C_{60} for X = He, Be, F^-, Br^- are shown in Figure 2. With higher mass of the central atom, the IR peak occurs in smaller frequency but the peak intensity also decreases. One important aspect of C_{60} is its tightly bound structure which remains intact in the encapsulated form. Thus X@C_{60} is both structurally strong and shows IR activity in the terahertz frequencies, though the intensity is small due to the weak coupling of the central atom to the cage.

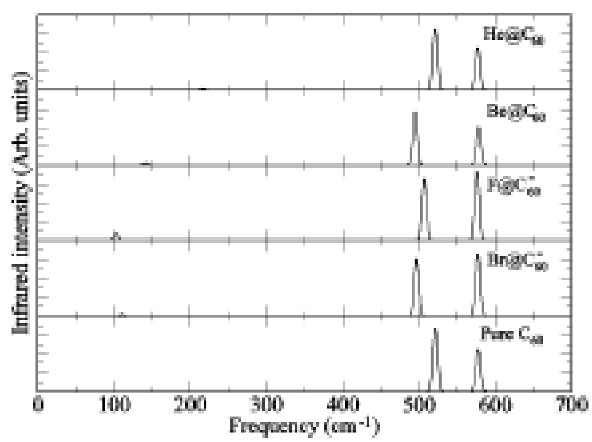

Figure 2: The infrared spectra of the X@C_{60} clusters in the range 0-700 cm^{-1}. The spectrum of pure C_{60} is given for reference.

References:

1) M. R. Pederson and K. A. Jackson, Phys. Rev. B. 41, 7453 (1990); K. A. Jackson and M. R. Pederson, Phys. Rev. B. 42, 3276 (1990); and M. R. Pederson and K. A. Jackson, Phys. Rev. B. 43, 7312 (1991).
2) D. V. Porezag and M. R. Pederson, Phys. Rev. A 60, 9566 (1999).
3) D. V. Porezag and M. R. Pederson, Phys. Rev. B. 54, 7830 (1996).

Author and Contact: Mark R. Pederson
Organization: Center for Computational Materials Science, Naval Research Laboratory, Washington, DC, 20375-5345
Author: Tunna Baruah
Organization: Department of Physics, State University of New York at Stony Brook, Stony Brook, NY, 11794
Resources: IBM SP3 at MHPCC and Linux cluster at Center for Computational Materials Science, Naval Research Laboratory
Sponsorship: Office of Naval Research Grant No. N000140211046, NSF Grant NIRT-0304122, and the DoD High Performance Computing CHSSI Program and Challenge Program

Hawaii Ferry Performance Analysis

Robert Dant, D. J. Fabozzi, J. Bergquist

AFRL/MHPCC has established a server environment under Project ENDEAVOR (Environment for Design of Advanced Marine Vehicles and Operations Research), consisting of three hardware servers powering a suite of commercial, open-source, and custom-developed software servers. AFRL/MHPCC (in partnership with SAIC, Navatek Ltd, and the University of Hawaii) established a capability to store Advanced Marine Vehicle (AMV) performance information to a database management system, and also established a GIS system to visualize world ocean environment data on the Web, including wind speed, current, wave, and bathymetry data. These resources were made available to support the efforts of the National Defense Center of Excellence for Research in Ocean Sciences (CEROS) in analyzing the performance characteristics of an advanced marine vehicle operating between the Hawaiian Islands.

Research Objectives: The goal of this research activity was to use the ENDEAVOR software framework, including the Large Amplitude Motions Program (LAMP) to perform nonlinear motion analyses to evaluate sea-keeping, stability, and operational ride characteristics of an advanced high-speed catamaran vessel (passenger and car ferry) operating on specified transit routes in littoral waters around the Hawaiian Islands.

Methodology: The methodology used was to: 1) investigate criteria for motion stability and ride characteristics for an inter-island ferry model of a predetermined size and shape (Figure 1), 2) integrate LAMP ship simulation code with an ocean wave model and GIS data (NOAA's Wave Watch III data), 3) generate transit waypoints along a given ferry route and determine sea-state conditions at these waypoints, and 4) analyze the performance characteristics of the high-speed ferry model by generating motion results.

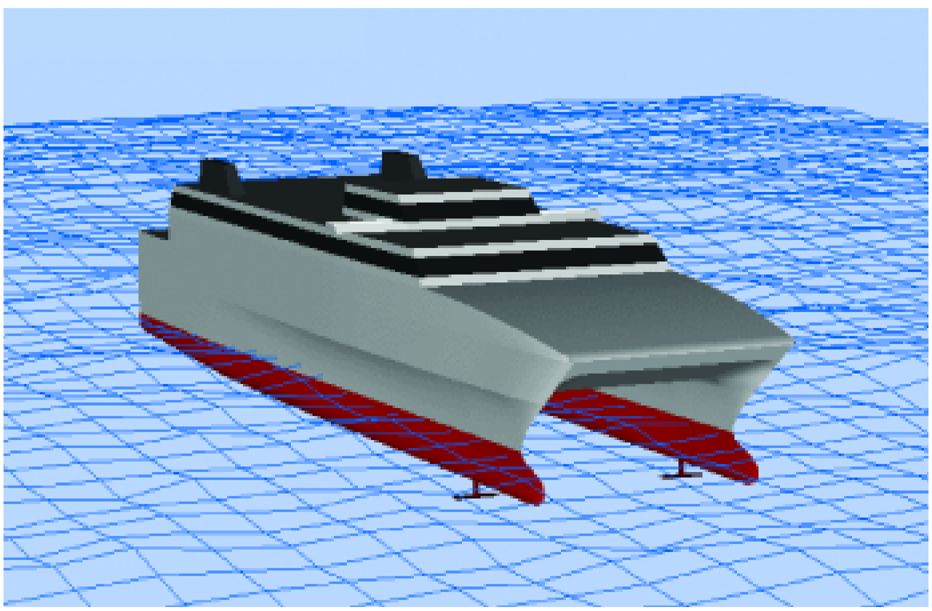

Figure 1. Generic ferry model used for the analysis of ship performance between islands for various routes and sea-state conditions. Overall ferry length ~300 ft., overall width ~86 ft., draft ~14 ft., and hull width ~19 ft.

AFRL/MHPCC hosted the LAMP code and provided technical support, as well as computational resources. Waypoints were determined at 30-minute increments using the ENDEAVOR environmental model and a specified time of year to determine the sea-state conditions at each waypoint along specified routes. AFRL/MHPCC also provided technical assistance in ocean wave database management, determination of routes using GIS, computation of near-shore wave conditions along specified routes using physics-based models, and storage of near-shore wave data.

Results: A capability was successfully demonstrated for route selection and operational planning of ferry and other marine vehicles (naval and commercial) operations. The technical and performance goals were achieved. The study showed that a typical wave spectrum does not reflect the true ocean wave conditions in the region of interest, and identified the need to consider region- and route-specific ocean environment for a high-speed ferry operation in different seasons. The study also demonstrated the capability of the ENDEAVOR system to provide technical capabilities and services for naval and commercial vehicle operations in the Hawaiian Islands.

A dynamic simulation was developed that evaluated the motions of a generic catamaran ferry model along two routes between Honolulu and Kahului (Maui), passing to opposite sides of Molokai (Figure 2). Wave conditions were determined using a University of Hawaii (UH) Wave Model, based on historical records for 7 October 2002, with 30 minute intervals from 0800 to 2000 hours, using approximately three hours transit time at 35 knots. A ship motion history along each route was evaluated using six 30-minute LAMP time-domain simulations. Peak response statistics were computed from predicted motion responses.

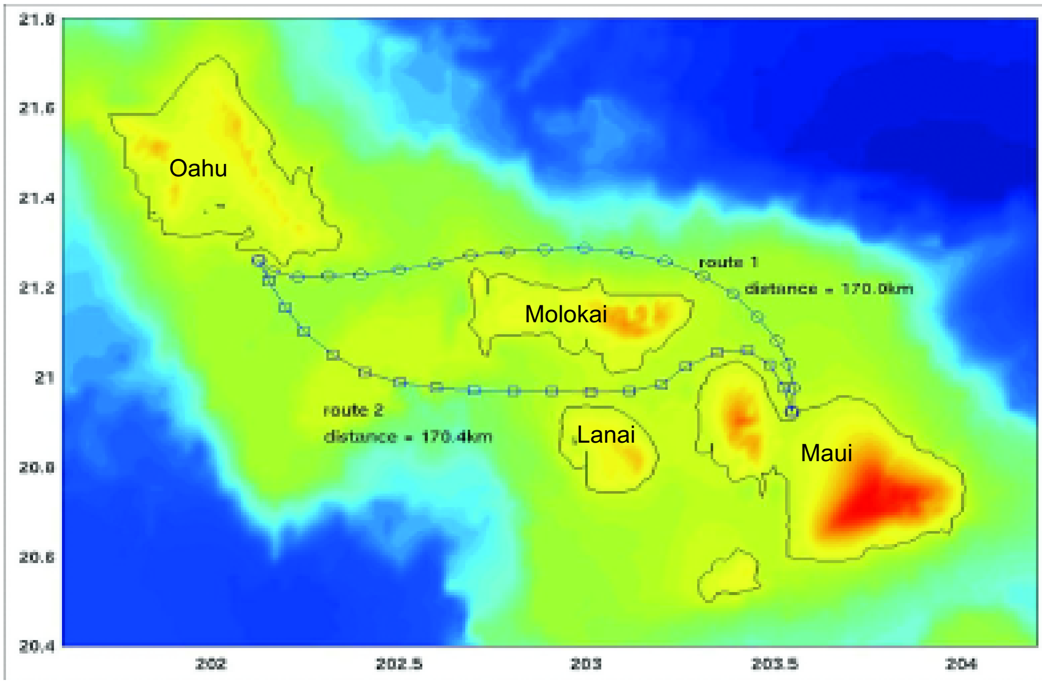

Figure 2. Simulated high-speed ferry routes were analyzed between Hawaiian Islands based on a ~3 hour transit at 35 knots. Route 1 is from Honolulu (Oahu) to Kahului (Maui) and transits north of the island of Molokai. Route 2 transits south of Molokai and passes through an extensive Whale Sanctuary area. Twenty waypoint locations were chosen along each route, with 30-minute intervals from 0800 to 2000 hours.

Significance: This research project was successful in conducting a scientific evaluation of a high-speed passenger and car ferry performance in a region-specific (Hawaiian Islands) ocean environment for differing seasonal conditions. The project also demonstrated the capability of route selection and operational planning for ferry and other marine vehicle (naval and commercial) operations. This information is now available for future design considerations of marine vehicles to operate around the Hawaiian Islands.

Author and Contact: Robert Dant
Author: D.J. Fabbozi
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Maui, HI, 96753
Author: J. Bergquist
Organization: SAIC Ship Technology Division, 134 Holiday Court, Suite 318, Annapolis, MD, 21401

Resources: Dell PowerEdge 2650 Dual Processors, ENDEAVOR PC Server, Dell Desktop; compute nodes - 128 nodes (256 processor), and *Huinalu* Linux cluster system running RedHat 9 operating system at MHPCC

Sponsorship: Air Force Research Laboratory

High-Resolution Forecasts to Support AMOS Using MM5

Kevin Roe and Duane Stevens

The Hawaiian Islands contain a variety of microclimates in a very small region. Some islands have rainforests within a few miles of deserts; some have 10,000 foot summits only a few miles away from the coastline. Because of this, weather models must be run at a much finer resolution to accurately predict in these regions. NCAR's Mesoscale Model Version 5 (MM5) is run from a coarse 27 km resolution (surrounding an area of approximately 5000 by 5000 km) nested down to a 1 km resolution daily. Since the computational requirements are high to accomplish this in a reasonable time frame (as to still be a forecast) MM5 is run in parallel on MHPCC's IBM SP4s. Utilizing 32 processors the MM5 model is run daily over the above conditions in approximately six hours. These forecasts have been in place for over two years now and are being utilized by operators at the telescope on Haleakala, Maui.

Research Objectives: The telescope operations on Haleakala are highly dependent on weather conditions on the Hawaiian Island of Maui. If the wind speed is too high then the telescope cannot be utilized. Problems also exist if there are clouds overhead. Rainfall and relative humidity are also a factor in determining the capabilities of the telescopes. In order to effectively schedule telescope operations, an accurate weather prediction is extremely valuable. Current forecasts that are available from the National Weather Service (NWS) give good indications of approaching storm fronts but only at the coarse level (30-50 km resolution). Because of this and the location of the telescope on Maui this can be insufficient for their needs. The additional benefit of the telescope operators having access to an accurate forecast (even for only a day in advance) is that they can still perform some scheduling. If a storm is predicted they can plan maintenance for this time period. This allows them to function more effectively by giving them the capability to schedule downtime. This in turn saves time, improves operating efficiency, and potentially saves money.

Daily Operations: Every night at midnight Hawaiian Standard Time (HST), a PERL script is run to handle all the operations necessary to produce a forecast, prepare the data, and post it to the MHPCC Web page (http://weather.mhpcc.edu/mm5). The procedures the script executes are:

1) Determine and download the latest global analysis files from NCEP for a 48-hour simulation,
2) Begin processing by sending these files through MM5's REGRID program,
3) Take the output data files from REGRID and input into INTERPF,
4) Prepare the MM5 model for the current simulation,
5) Submit the MM5 run to MHPCC's IBM SP4 (*Tempest*) for execution (daily reservation starting at 1 A.M.),
6) Average daily run requires 5-6 hours for completion on 32 processors (a single P4 node),
7) Data is output in one-hour increments,
8) Data is processed in parallel to create useful images for meteorological examination,
9) Convert images to a Web viewable format,
10) Create the Web pages for these images, and
11) Post Web pages and images to MHPCC's Web site.

Figure 1. Maui Space Surveillance Site located atop Mt. Haleakala.

Web Output: Now that the above processes have created images, they must be made available for the telescope operators. This is accomplished by posting the images to the MHPCC Web page; specifically, http://weather.mhpcc.edu/mm5. This title page gives the user the option of what area and resolution they would like to examine. From the title page, the user can select the all island area at a 27 or 9 km resolution, one of the four counties (Hawaii, Maui, Oahu, and Kauai) at a 3 km resolution, or the summit of Haleakala at a 1 km resolution. Once one of the above has been selected, the user is transported to a web page that initially includes an image of the temperature in the selected area. On the regional Web page, the viewer can select to see the previous or next image through the use of a small JavaScript. If the viewer prefers, an animation of the images (in one hour increments) can be started and stopped. Finally, the user can select any of the images from a pull-down menu. If the viewer would like to change the field being examined, a pull down menu on the left side of the page will transport the user back to the main menu or allow them to choose a different field. Lastly, if the JavaScript becomes a problem for the viewer's browser, they have the option of being switched to a non-JavaScript equivalent version Web page.

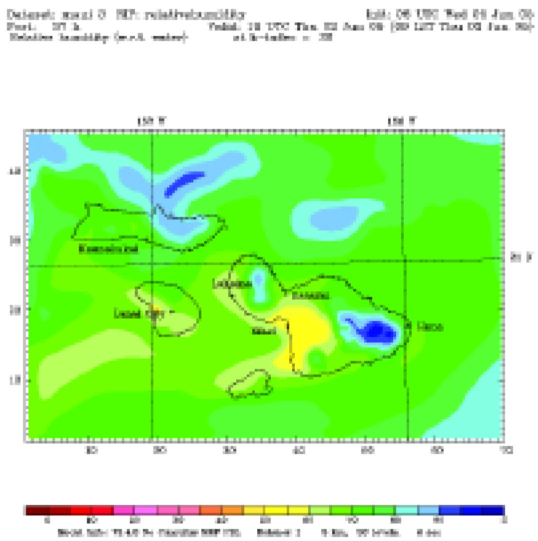

Figure 2. Website screen capture of weather modeling of the Hawaiian Islands.

References:

1) J. Michalakes, "The Same-Source Parallel MM5," *Journal of Scientific Programming*, 8, (2000): 5-12.
2) J. Michalakes, S. Chen, J. Dudhia, L. Hart, J. Klemp, J. Middlecoff, and W. Skamarock, "Development of a Next Generation Regional Weather Research and Forecast Model," Developments in Teracomputing, Proceedings of the Ninth ECMWF Workshop on the Use of High Performance Computing in Meteorology, Eds., Walter Zwieflhofer and Norbert Kreitz, World Scientific, Singapore, (2000): 269-276.
3) J. Dudhia, "A Nonhydrostatic Version of the Penn State/NCAR Mesoscale Model, Validation Test and Simulation of an Atlantic Cyclone and Cold Front," Mon. Wea. Rev., 121, (1993): 1493-1513.
4) Yi-Leng Chen and J. Feng, "Numerical Simulation of Airflow and Cloud Distributions over the Windward Side of the Island of Hawaii, Part I: The Effects of Trade Wind Inversion," American Meteorological Society, (May 2001): 1117-1134.
5) Yi-Leng Chen and J. Feng, "Numerical Simulation of Airflow and Cloud Distributions over the Windward Side of the Island of Hawaii, Part II, Nocturnal Flow Regime," American Meteorological Society, (May 2001): 1135-1147.
6) K. P. Roe and D. Stevens, "High Resolution Weather Modeling in the State of Hawaii," The 11th PSU/NCAR Mesoscale Model Users' Workshop, Boulder, C0, 2001.
7) K. P. Roe and D. Stevens, "High-Resolution Weather Modeling to aid AMOS," DoD HPCMP User Group Conference, Bellevue, WA, 2003.
8) K. P. Roe and M. Waterson, "High Resolution Numerical Weather Forecasting to Aid AMOS," 2003 AMOS Technical Conference, Wailea, HI, 2003.

Author and Contact: Kevin Roe
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Maui, HI, 96753
Author: Duane Stevens
Organization: Department of Meteorology, University of Hawaii at Manoa, 2525 Correa Road, HIG 350, Honolulu, HI, 96822
URL: http://weather.mhpcc.edu/mm5

Resources: IBM SP4 at MHPCC

Sponsorship: Air Force Research Laboratory

Hawaii Fire Danger Rating System

Kevin P. Roe and Francis M. Fujioka

Forest lands in Hawaii provide clean water and air for 1.2 million people, and habitat to rare plants and animals found nowhere else in the world. To help natural resource managers protect Hawaii's forests, the USDA Forest Service and MHPCC recently installed a new Hawaii Fire Danger Rating System at MHPCC, which combines the MM5 high resolution weather modeling framework with the National Fire Danger Rating System (NFDRS). The project built on previous work created terrain and fuels map databases necessary to evaluate fire danger.

Project Objectives: Upon completion, the Hawaii Fire Danger Rating System will incorporate high resolution weather and fire danger modeling subsystems to inform land managers of fire risks in Hawaiian forests and wildlands (Figure 1).[1] The System accounts for weather, fuels, and terrain characteristics in predicting fire ignition, spread, and heat energy release potential within a 48-72 hour forecast period. Wildland fire protection agencies will be able to create customized weather and fire danger maps on demand using an interactive web-based map service.

Figure 1. Wildfires in Hawaii typically burn through fine fuels such as the fuelbed shown here.

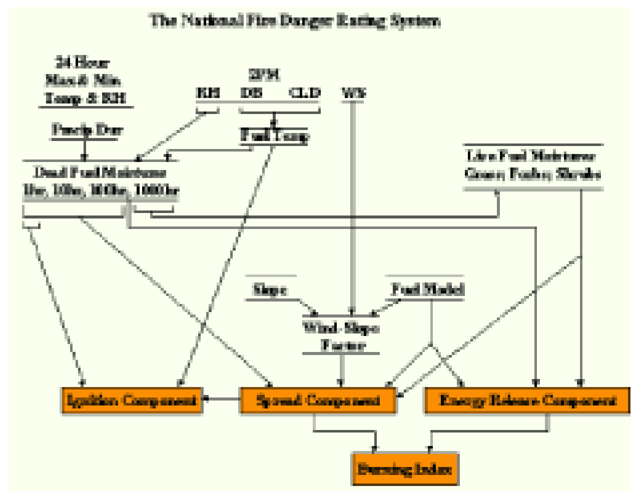

Figure 2. The Hawaii Fire Danger Rating System follows the design of the National Fire Danger Rating System. Ignition, spread, and energy release characteristics of a fire are determined from input streams of weather, fuel, and terrain data that describe the fire environment.

Methodology: The daily process of producing fire danger rating forecasts begins with MM5 model runs at MHPCC that produce fire weather forecasts for the counties of Kauai, Oahu, Maui, and Hawaii, at a grid interval of 3 km.[2] A fire danger rating post-processor uses the MM5 forecasts to calculate corresponding fire danger rating variables, integrating the additional effects of fuel and terrain conditions, which are assumed to be static. The weather data-relative humidity (RH), temperature (DB), cloudiness (CLD), and precipitation (Precip Dur) determine the moisture content of the dead fuel elements in the fuels map database, hence the ignition potential (Figure 2). An underlying fire behavior model within the NFDRS incorporates wind (WS) and topographic slope effects to estimate fire spread and energy release potentials. The Burning Index, derived from the Spread and Energy Release Components, is theoretically proportional to the expected flame lengths, given the fire environment inputs.

The fire danger rating components quantify the dimensions of the expected fire potential, thus providing fire protection agencies with information to respond with appropriate prevention, dispatching, and, if necessary, resource pre-positioning strategies. However, each agency may have its own critical thresholds of fire danger rating components that trigger management actions. This requires flexibility in extracting information from the danger rating system.

The Hawaii Fire Danger Rating System meets this requirement with a web-based geographic information system (GIS). The user can create weather and fire danger index maps for any time in the forecast period, and at the desired map scale (Figure 3). Other features such as roads, agency boundaries, place names, and even color satellite images may be selectively displayed, thus giving the user a wide variety of information to determine an appropriate management action.

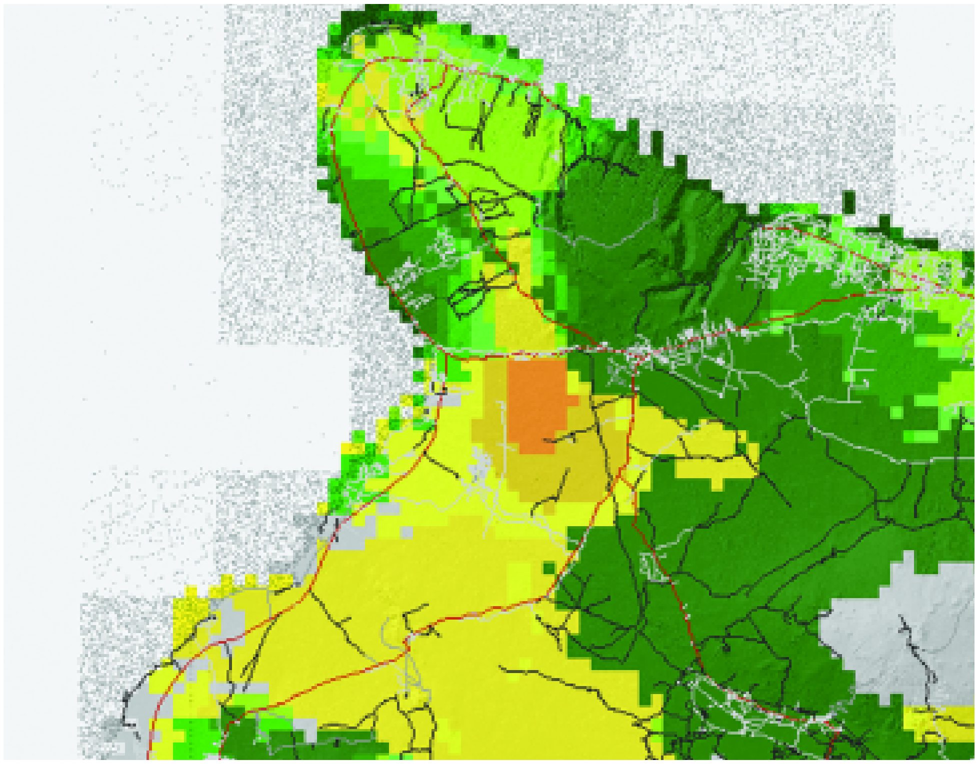

Figure 3. Forecast map of Energy Release Component for the North Kohala area of the Big Island. This map was created from the GIS Web server for the HFDRS.

Future Work: The Hawaii Fire Danger Rating System informs fire managers of potential fire problems. Once an incident occurs, good predictions of fire spread are needed, which requires even higher resolution weather, fuels, and topography. Future work will attempt to drive model grid intervals to less than one kilometer, and integrate the weather model with a fire behavior model.[3]

References:

1) F. M. Fujioka, D. R. Weise, and R. E. Burgan, "A High Resolution Fire Danger Rating System for Hawaii," Preprints, 3rd Symposium on Fire and Forest Meteorology, 9-14 January 2000, Long Beach, CA, American Meteorological Society, 103-106.
2) K. P. Roe and D. Stevens, "High Resolution Weather Modeling in the State of Hawaii," 11th PSU/NCAR Mesoscale Model Users' Workshop, Boulder, CO, 2001.
3) F. M. Fujioka, "PEGASUS: An Integrated Weather/Wildfire Modeling System," Preprints, 2nd Symposium on Fire and Forest Meteorology, 11-16 January 1998, Phoenix, AZ, American Meteorological Society, 38-41.

Author and Contact: Kevin P. Roe
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Author: Francis M. Fujioka
Organization: USDA Forest Service, 4955 Canyon Crest Drive, Riverside, CA, 92507

Resources: IBM SP4s at MHPCC

Acknowledgement: The accomplishments of this project would not have been possible without the invaluable assistance of Robert E. Burgan (retired research forester, USDA Forest Service), John W. Benoit (Computer Specialist, USDA Forest Service, Riverside, CA), and Andrew E. Wilson (GIS specialist, USDA Forest Service, Portland, OR).

Maritime Synthetic Range (MSR)

Carl Holmberg, Robert Dant, Thomas Meyer, Aaron Steigerwald

The Air Force Research Laboratory Maui High Performance Computing Center (AFRL/MHPCC), under the auspices of the Office of Naval Research (ONR), is pursuing the Maritime Synthetic Range (MSR) program to provide more opportunities for realistic tactical training of sailors aboard ship, aloft, and ashore by linking their combat stations to simulators and computerized models over secure networks (Figure 1), such as SDREN (Secret Defense Research and Engineering Network). Frequent and varied training exercises of arbitrary size (including widely disbursed units) will be possible, limited only by available network connections and bandwidth. At the same time, the Fleet will realize considerable savings in organizational effort and costs, as compared to traditional at-sea exercises.

Research Objective: The initial MSR program objectives at AFRL/MHPCC were: 1) to create a prototype virtual range capability (Figure 2), with simulated targets and physical environments hosted on dedicated high performance computers at AFRL/MHPCC, and linked to live and simulated naval assets based at the Pacific Missile Range Facility (PMRF), and 2) to provide remote 2- and 3-D event visualizations within the virtual range (Figure 3).

Figure 1. By feeding realistic data to their tactical consoles, naval units can practice engagements with each other as well as against computer simulated opponents, free from geographic constraints, and linked by secure networks to land-based, high-speed computers.

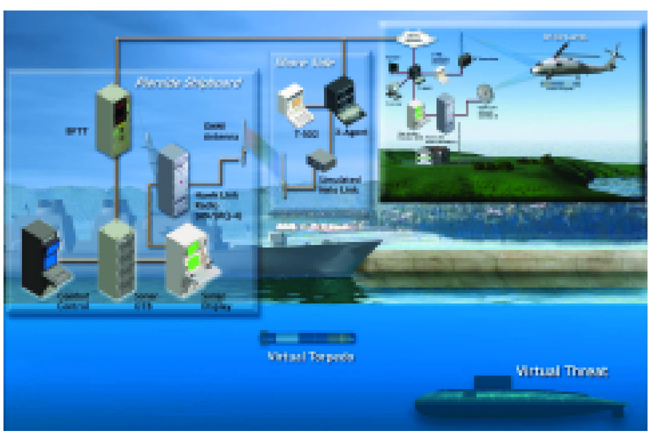

Figure 2. Live, virtual, and constructive elements are tied together to create various Maritime Synthetic Range (MSR) scenarios.

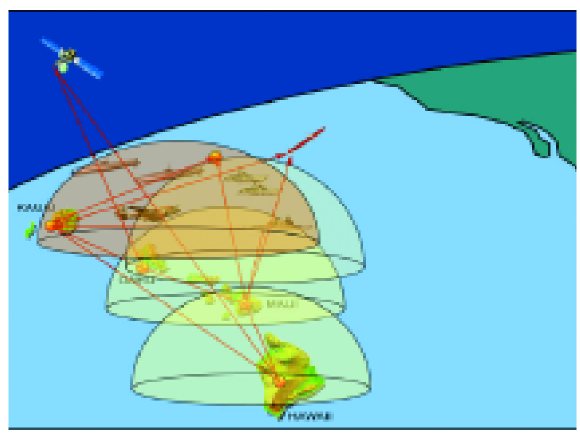

Figure 3. Various 2- and 3-D event visualizations are created with a "virtual range" environment.

Methodology: Research efforts for the MSR program (Year 1) at AFRL-MHPCC focused on: 1) installing dedicated secure network and computing resources, 2) installing Joint Semi-Automated Forces (JSAF) and Simulation Display (SIMDIS) software (Figures 4 and 5), 3) training the staff engineers in the use and modification of JSAF, and 4) developing gateway software to enable JSAF and SIMDIS to interact. This gateway provided users with enhanced visualization of JSAF scenarios in real-time. SIMDIS, developed by the Naval Research Laboratory (NRL), was selected as the MSR platform for high-fidelity analysis and display of test and training mission data.

Results: AFRL/MHPCC purchased and installed a 32-processor Linux cluster system for exclusive MSR use, a stand-alone MSR server/workstation, and secure network communication equipment. JSAF tactical modeling and SIMDIS visualization software were installed on the workstation and JSAF scene-generation capabilities were developed based on NRL's SIMDIS synthetic display environments. Gateway software was developed to translate messages between these two packages and used in a NAVAIR simulated, tactical-engagement demonstration at the 2004 Interservice/Industry Training, Simulation and Education Conference (I/ITSEC).

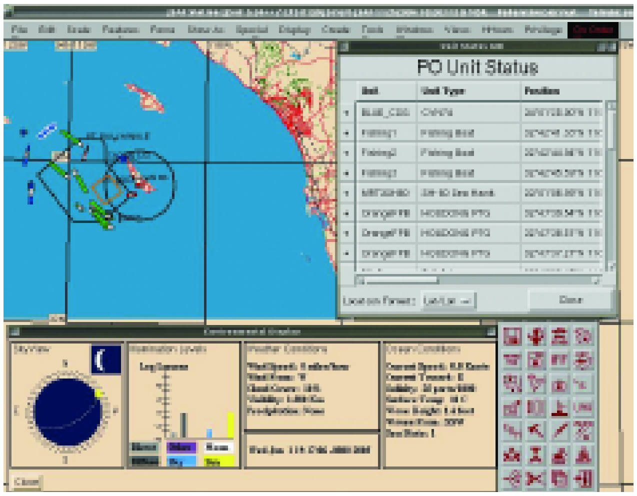

Figure 4. JSAF is the core of the simulated range system for MSR.

Figure 5. SIMDIS was developed by NRL and is a key operational display system for MSR.

Significance: A key to the U.S. warfighters' advantage in the field is frequent, effective training as a unit. The most effective training occurs within a realistic combat environment, which has traditionally led to exercises in the field. However, field exercises are expensive in terms of the time, personnel, assets, and funding they consume that might otherwise have gone to support actual deployments. AFRL/MHPCC is laying the groundwork for future advances in WAN-enabled virtual exercises. As the MSR infrastructure becomes reality and reaches out into the Fleet and Joint Forces, combat units will enjoy more opportunities to maintain their edge without straining existing training and operations resources.

Author and Contact: Carl Holmberg
Authors: Robert Dant, Thomas Meyer, Aaron Steigerwald
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Resources: IBM Linux Cluster (32 Xeon Pentium IV processors, 3.06 GHz) and one Dell Precision 360 MiniTower Server/Workstation at MHPCC
Sponsorship: Office of Naval Research

Airborne Laser Atmospheric Decision Aid (ADA) Testing

Frank H. Ruggiero and Kevin P. Roe

The Airborne Laser (ABL) is being developed as one element of our nation's Ballistic Missile Defense System. The ABL Element Office is developing an Atmospheric Decision Aid (ADA) to diagnose and forecast the location and magnitude of optical turbulence in the upper troposphere and lower stratosphere. The ADA will use as input real-time numerical weather prediction (NWP) forecasts provided by the Air Force Weather Agency (AFWA). This project evaluated the overall performance of the current operational and next generation NWP models and whether the current model grid spacing used in the AFWA operational models is sufficient for the optical turbulence parameterizations used in the ADA. In addition, an updated optical turbulence parameterization was compared with original parameterization put in the ADA. Results show that the AFWA's new NWP model run at planned grid spacing will provide sufficient input for both the current and new optical turbulence parameterization. The updated optical turbulence parameterization showed significant improvement over the existing parameterizations.

Background and Objective: The Airborne Laser (ABL), an element of our nation's Ballistic Missile Defense System (BMDS), will detect and destroy boosting ballistic missiles. The ABL Element Office is managing the development of an Atmospheric Decision Aid (ADA) to diagnose and forecast the location and magnitude of optical turbulence in the upper troposphere and lower stratosphere. Optical turbulence is the fluctuation of density in the atmosphere and acts to defocus laser beams and thus reduce their effective range. Layers of intense optical turbulence, while possibly stretching hundreds of kilometers in the horizontal, tend to occur on vertical scales on the order of 100 meters. Mission planners responsible for the ABL will need to know the areas of strong optical turbulence fields in order to optimally locate the ABL's orbit to protect the asset while maximizing the lethality capability. In order to provide mission planners with real-time forecasts of optical turbulence, the ABL Element Office is developing an Atmospheric Decision Aid (ADA) to diagnose and forecast the location and magnitude of optical turbulence in the upper troposphere and lower stratosphere. The ADA uses output of a numerical weather prediction (NWP) model and runs a postprocessor that contains a turbulence parameterization that quantifies optical turbulence from weather variables.

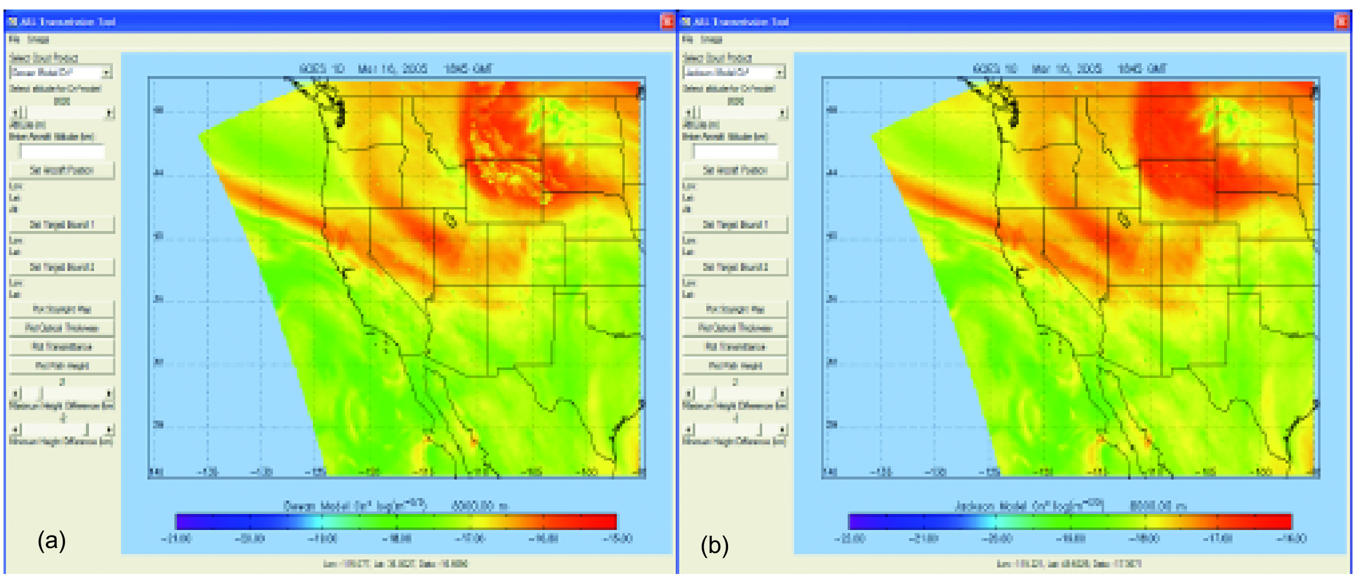

Figure 1. Predictions of the refractive index structure constant (C_n^2), a measure of optical turbulence, at a constant height surface of 8 km AGL for (a) the Dewan *et al.* (1993) and (b) the Jackson (2004) optical turbulence parameterizations displayed on a prototype ABL ADA. (Figure courtesy of Maj. J. D. Cetola, Airborne Laser Element Office)

The Air Force Weather Agency (AFWA) will send its real time operational mesoscale model forecasts to the ADA. Currently, the AFWA uses the Fifth Generation Penn State/National Center for Atmospheric Research Mesoscale Model[2] (MM5; Grell *et al.* 1994,[1] Michalakes 2000) as its operational model. It will soon transition to the new Weather Research and Forecast[3] (WRF; Skamarock *et al.* 2005) Model for its operations. The objective of the work of this project is two fold. First, we want to determine if the WRF model run at AFWA's operational grid spacing (both horizontally and vertically) will be sufficient for input to the turbulence parameterizations. The second objective is to produce verification statistics on the Dewan *et al.* (1993)[4] and Jackson (2004)[5] optical turbulence parameterizations.

Results: In 2004-2005 we conducted verification runs of MM5 and WRF NWP models combined with both the Dewan *et al.* (1993) and Jackson (2004) optical turbulence parameterizations on the MHPCC IBM SP/RS6000 SP computer for two different field experiments where optical turbulence measurements were available. MM5 and WRF were run with varying configurations of grid spacing. Results showed the WRF model running with the current operational grid spacing should be sufficient input to the optical turbulence parameterizations. Comparison of results of the Dewan *et al.* (1993) and Jackson (2004) optical turbulence parameterizations show that the Jackson optical turbulence parameterization is superior to Dewan. However, the two parameterizations do show complementary behavior and could be useful to run together as a two member ensemble.

References:

1) G. A. Grell, J. Duhia, and D. R. Stauffer, "A Description of the Fifth Generation Penn State/NCAR Mesoscale Model (MM5)," NCAR Tech. Note T/N-398 + STR, (1995); 122 pp, [Available from NCAR Information Services, P.O. Box 3000, Boulder, CO, 80307, (http://www.mmm.ucar.edu/mm5/documents/mm5-desc-doc.html)].
2) J. Michalakes, "The Same-Source Parallel MM5, Sci. Programming," 8, (2000): 5-12. (http://www.mcs.anl.gov/~michalakes/JOSP-michalakes.ps)
3) W. C. Skamarock, J. B. Klemp, J. Dudhia, D. O. Gill, D. M. Barker, W. Wang, J. G. Powers, "A Description of the Advance Research WRF Version 2", NCAR Tech. Note T/N - 468 + STR, (2005); 100 pp, [Available from NCAR Information Services, P.O. Box 3000, Boulder, CO, 80307] (http://www.mmm.ucar.edu/wrf/users/pub-doc.html)].
4) E. M. Dewan, R. E. Good, B. Beland, and J. Brown, "A Model for Cn^2 (Optical Turbulence) Profiles using Radiosonde Data," Phillips Laboratory Technical Report, PL-TR-93-2043, (1993), ADA 279399.
5) A. Jackson, "Modified-Dewan Optical Turbulence Parameterizations," Air Force Research Laboratory Technical Report AFRL-VS-HA-TR-2004-1116, 2004, ADA 432901.

Author and Contact: Frank H. Ruggiero
Organization: AFRL, Space Vehicles Directorate, Hanscom AFB, MA, 01731-3010
Author: Kevin P. Roe
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, Hi, 96753

Resources: IBM SP3/SP4 at MHPCC

Sponsorship: Air Force Research Laboratory and ABL Program Element Office

High Performance Computing and Visualization Support for Project Albert: Analysis of Simulations of Combat Operations and Operations Other Than War

Bob Swanson, Michael Coulman, Ron Viloria

The Marine Corps Warfighting Laboratory's (MCWL) Project Albert involves research to assess the general applicability of the "New Sciences" to combat operations and operations other than war. While simulations based on these New Sciences (or complexity theory, which models behavior and interaction at the entity level) are considered part of an infant science, they do provide insight into evolving patterns of macroscopic behavior that result from collective interactions of individual agents. Simulations based on entity-level interactions represent a significantly different approach from the traditional attrition estimation techniques based on Lanchester equations, which assert that the loss rate of forces on one side of a battle is proportional to the number of forces on the other side. MCWL's Project Albert is an effort to investigate how complexity theory may be applied to combat in a manner to augment and, perhaps in some cases, replace Lanchester modeling in the future.

Research Objectives: MCWL is continuing the development of a complex systems analyst's toolbox for exploiting emergent, collective patterns of behavior on the battlefield, and is currently using several multi-agent-based simulations of notional combat. Current models that are being employed include: Map Aware Non-uniform Automat (MANA), Pythagoras, Simulation Of Cooperative Realistic Autonomous Transparent Entities (SOCRATES), PAX, NetLogo, and the Irreducible Semi-Autonomous Adaptive Combat (ISAAC) model. These models simulate the interaction between two or more variable size forces of agents. The action of each agent is determined by a small rule set specified by parameters, such as the agent's ability to sense its surroundings and to communicate with other agents. Besides the common physical parameters measured by the model (such as range of fire and probability of kill), more abstract concepts are also modeled. Such concepts can include an agent's attraction to friendly and opposing forces and the influence of behaviors based on determinable thresholds, such as the tendency of an agent to follow orders. The magnitude and granularity of these independent variables provide the analyst with great flexibility in simulating various hypotheses. However, this same flexibility, coupled with the stochastic nature of the simulations, requires a significant computational capability to determine likely outcomes with any statistical significance. This requirement becomes even greater for the analyst who wants to study hypotheses over multiple varying independent parameters, a requirement that can easily overtax the capability of a single personal computer.

Methodology: The Maui High Performance Computing Center (MHPCC) Project Albert contractor team has performed computer systems and software engineering services in concert with military analysts from MCWL and other supporting contractors, to develop methodologies and tools needed for the large-scale analysis of agent-based distillation models. This development includes tools to define a problem set, request a job submittal, assist MHPCC personnel in running and managing the jobs, marshal the sometimes massive output of the jobs, and tools to help the end-user analyst visualize the numeric results created by these chaotic systems. The models are largely cross-platform in nature, but MHPCC has ported those models, where possible, to work on MHPCC's IBM Linux cluster, the IBM SP parallel supercomputer, and a cluster of Windows workstations unique and dedicated to Project Albert. The primary output of the many model runs is a set of numeric measures of battlefield effectiveness. In addition, tools produced at MHPCC can extract time/space results from certain models that add movie-like visualization capabilities of the battlefield state. The information generated from this exhaustive execution of model simulations, along with the associated mechanisms for visualization, will provide analysts with the tools required to investigate multiple hypotheses with statistical significance.

Figure 1. Three Dimensional Scatter Plot.

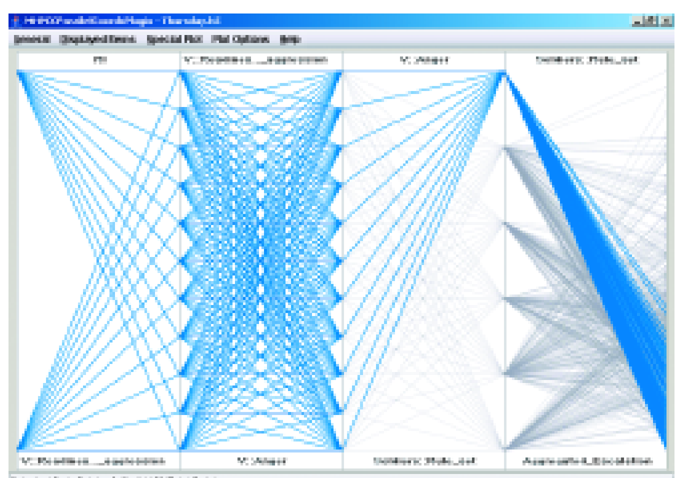

Figure 2. Parallel Coordinate Plot.

Results: The MHPCC contractor team created a new Visualization system for Project Albert called "Albert Visualization Abstraction Toolkit for Analysis and Research (AVATAR)." This tool was made available to researchers in 2004, but has since been extensively enhanced.

The AVATAR system is designed with a core "Data Console" that allows the user to select the portions of the data space they wish to view. It should be noted that a large set of runs from Project Albert can create millions of data points. The Console makes it easier for an analyst to grasp a large set of numeric information.

In order to view the selected information, the AVATAR system supplies a simple "plugin" interface. Thus, any Java programmer can produce a plotting plugin that views and/or analyzes the Measures of Effectiveness (MOE) output values from Project Albert data farming runs. Recent work on AVATAR has included the creation of several new data display plugins.

Figure 1 shows a three-dimensional scatter plot. The user can easily "navigate" through this data space, using their mouse and/or keys to rotate, translate, and zoom the scene.

Figure 2 shows a Parallel Coordinate plot. This type of plot is very useful when the result space contains a high number of data dimensions, a very common situation seen in Project Albert data farming output results. One of the strengths of parallel coordinate plots, is the ability of the user to "brush", or highlight, ranges of data points, in order to determine the relationship between their component values.

Figure 3 shows a "Box and Whiskers" plot. Each display column represents statistical information about that data point. The "box" is bounded by the 1st and 3rd quartiles of all the data values at that location. The "whiskers" are markers for the minimum and maximum values. The large black dot is the mean of all the values. This type of plot allows the user to gain an understanding of the "shape" of the data, and to further explore the data points by viewing additional distribution bar charts.

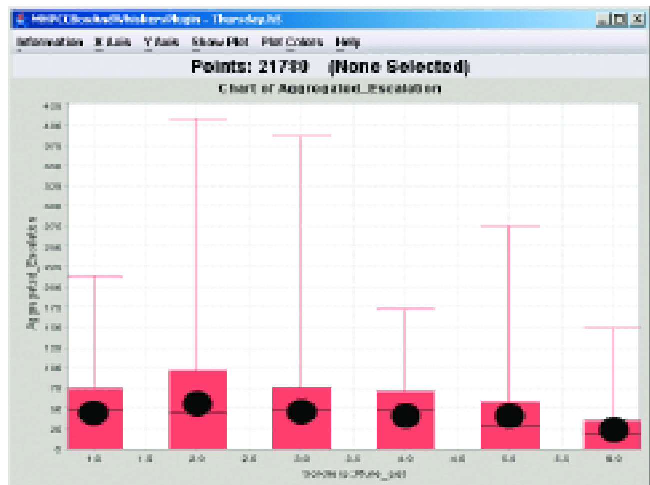

Figure 3. Box and Whiskers Plot.

Figure 4 shows a two-dimensional scatter plot with jittering added. When a plot is "jittered", small random variations are combined with the results data, causing each data point to occupy a slightly different location on the screen. Using this method, it is much easier for the user to determine the density of information at any location on the plot.

Work on this tool is ongoing, with feedback from real-world customers helping MHPCC to make the product more useful.

Significance: This initiative has led to a continuing evolution and refinement of the new generation of analytical models and tools for Department of Defense researchers. It is anticipated that MCWL will continue to require further development of simulation, visualization, and statistical methodologies that may enhance MCWL's ability to evaluate the applicability of these simulations to actual combat doctrine. These efforts will expand the existing infrastructure to incorporate new technologies.

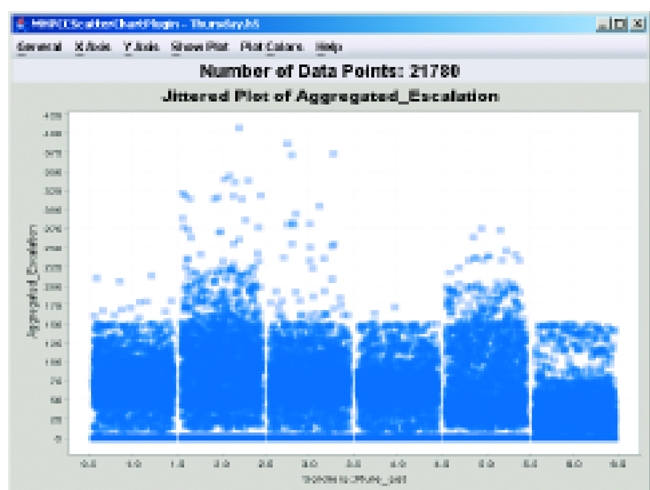

Figure 4. Two-Dimensional Scatter Plot.

Author and Contact: Bob Swanson
Authors: Michael Coulman and Ron Viloria
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Resources: IBM P3, IBM Linux Supercluster, and Windows clusters at MHPCC
Sponsorship: Marine Corps Warfighting Laboratory (MCWL)

Easy Enterprise Grid Computing

Michael S. Coulman

The Maui High Performance Computing Center (MHPCC) supports the Marine Corps Warfighting Laboratory (MCWL) Project Albert, whose goal is to develop tools and methods that bring combat simulation into the 21st century. MCWL uses agent-based combat simulations, which model both classical force strength attributes and behavioral attributes (morale, loyalty, bravery, obedience, etc.), to provide very complex combat simulations. This complexity, and the military analyst's desire to vary both multiple parameter values and the randomness of the simulation, quickly overwhelms the ability of a single CPU to deliver statistically significant results in an acceptable amount of time.

Research Objectives: Complex agent-based simulations may be run on a computational grid employing multiple CPUs to better deliver results in a time effective manner. To support multiple agent-based simulations, a generalized computational grid environment requires the following characteristics: 1) Extensibility, 2) Platform and operating system independence, 3) Ability to Scale to arbitrarily large jobs, 4) Optimal throughput, and 5) Dynamic resource allocation.

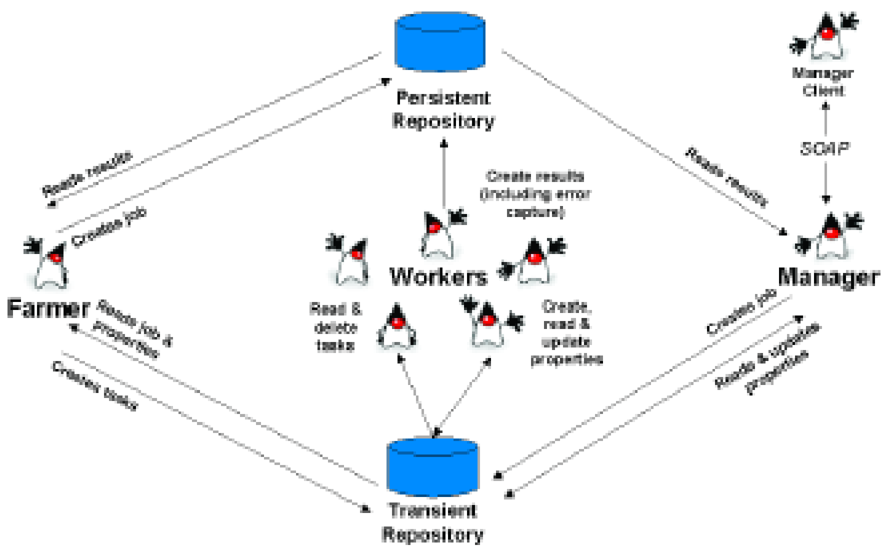

Figure 1. The Work Flow.

Since Project Albert supports multiple agent-based simulations on a variety of platforms, a boutique solution tailored to a specific simulation or platform would not be appropriate. In normal circumstances, Project Albert has access to approximately 50 CPUs at MHPCC. However, during International Workshops, Project Albert has access to some of MHPCC's larger resources, providing hundreds of CPUs. As a result, the grid must be able to rapidly change in size and scale. Most importantly, any computational grid must quickly respond to changing priorities so simulations are done *fast.*

Implementation: The analogy of a farm provides an intuitive conceptual understanding of our enterprise grid, called JEM (Job Execution Manager). All of the software entities, referred to as agents, that JEM uses are plain old Java objects (POJOs) that have direct analogs to roles in an agricultural work environment.

The Manager agent obtains seeds that the Farmer agent can use to plant crops, and furnishes the Farmer agent with a tool appropriate for planting various types of seeds. The Manager agent controls the allocation of Worker agents to particular crops, and furnishes them with the appropriate tool for performing the chores associated with each assigned crop. The Farmer agent is responsible for fetching the seeds that enable the planting of crops, and for dividing the planting into chores that will be performed by the Worker agents. Also, the Farmer agent is responsible for keeping records about the crops, and using those records to determine when the crop is ripe for harvest. The Worker agent does the chores. When the Worker agent completes a chore it sends the results to a persistent repository, and then gets another chore to do, until there are no chores left to do.

Neither the Farmer nor the Worker has a user interface or API. All interaction with them is via messages that they request from a transient repository, much like checking for voice mail or email. This uncoupled, highly cohesive architecture results in a robust, reliable set of services that are, by and large, immune to dependency related failure. For example, if a Worker should fail, not only do the other Workers keep on doing their chores, but the Farmer and Manager continue performing their responsibilities as well. The Farmer notices any incomplete chore resulting from Worker failure and resubmits it to the remaining Workers. The Manager notices the failed Worker, and posts alerts so it receives attention from outside of JEM.

The Manager Agent: The Manager is the user interface to the JEM farm. It provides a GUI that controls both the "how" and "when" for all aspects of JEM's operations. The Manager obtains "seeds" that the Farmer uses to plant crops and provides the tools the Farmer uses during planting. At any time the Manager may send Workers messages allocating their efforts to a specific crop's chores and provide the Workers with the tool to use to perform the chore. The Manager can also instruct Workers to stop performing chores, or to completely shut down.

Based on Worker assignments, performance metrics gathered by the assigned Workers, and the sequencing algorithm selected to order the chores, the Manager estimates when all the chores associated with each crop will be completed. The Manager monitors the health of each Worker, watching for abnormally long completion times as compared to its peers. The Manager also monitors the health of the JEM system, verifying that communication takes place as intended.

The Farmer Agent: The Farmer is responsible for "planting," the process of using the seeds the Manager provides containing information about the "crops." The Farmer deconstructs the crop planting into "chores," which are the independent, atomic units of work performed by Workers on the farm. The Farmer distributes the chores to a pool that the farm's Workers can draw from. During and after planting, the Farmer watches the crop and determines when it is "ripe," that is when the Workers have completed all the crop's chores. The Farmer can learn how to plant new crops on the fly by receiving additional "tools" from the Manager via the network.

The Farmer doesn't care which Worker does what chore, or how many chores any given Worker can do as compared to the others. The Farmer watches how many Workers are assigned to a given crop and how quickly the assigned Workers perform chores, on average, in order to scale chore distribution to ensure that no Worker is ever idle and no chore pool is ever larger than what seems reasonable to keep the assigned Workers continually busy. The Farmer also monitors Workers for errors, and upon discovery of an error the Farmer automatically resubmits that chore to the chore pool to be done again. The Farmer can replant a previously terminated, incomplete crop without duplicating correctly completed chores.

The Worker: The Worker, a common resource that may be shared by all the crops, is responsible for doing chores around the farm. JEM Workers are "smart" in two ways. First, they only attempt chores that they know they are able to do according to the platform and operating system that they run on. Second, Workers exclusively perform assigned chores until all assigned chores are exhausted, and then, if other chores are present and they have access to the tool to perform the chore, they will automatically perform those as well by random selection.

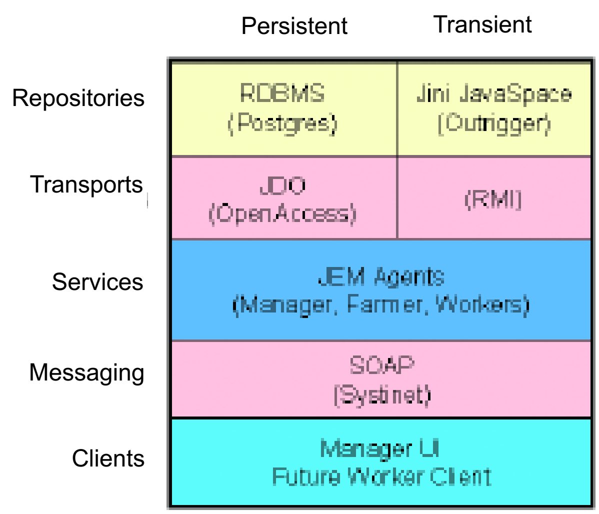

Figure 2. Architecture.

Also, like the Farmer, Workers learn to do new chores on the fly by getting additional tools from the network. If a Worker's CPU is needed for other tasks the Worker can sleep until the CPU is available and then resume doing chores upon command.

As previously stated, a loosely coupled, highly cohesive system is best able to provide a robust, reliable set of services that have no dependency related liabilities. In fact, JEM, as implemented, has *no* coupling between its agents. That is to say, the Manager has no knowledge of or reference to either the Farmer or any Worker. The same is true for both the Farmer and the Worker with respect to each other and the Manager. How then does a system with no coupling communicate? It communicates through messages sent via the network.

Transient Repository: A transient repository, a JINI JavaSpace, provides one network communication medium for JEM's agents. Messages about pending chores, system performance, agent status, shared data, and even bundled executable software objects, move from agent to agent via the transient repository.

Given the large magnitude of chores per crop with respect to the significantly smaller magnitude of Workers, the Farmer metes out chores at a measured rate so as to not overwhelm the system running the transient repository. Further, as Workers perform chores they provide performance metrics, via messages, that the Farmer uses to determine an interval to mete out additional chores. Thus, the system running the Farmer minimizes the number of CPU cycles spent checking to see if it is time to send additional chores to Workers via the transient repository.

Using this implementation, a 2 GHz Pentium 4 Mobile Sony Vaio notebook computer with 512 Mb of RAM running the JINI class server, the JINI JavaSpace, the Farmer, and the Manager, all in separate JVMs, farmed crops of up to 750,000 chores lasting four days to a pool of 50 Worker CPUs at MHPCC.

As previously stated, one or more chores comprise a crop. Additionally, every successfully completed chore produces computational output, referred to as a "result." The Java Data Objects (JDO) API provides a second network communication medium for JEM's agents using POJOs. The Farmer transmits top level information about the crop, and Workers transmit results to a persistent repository, in this case the Postgres ORDBMS, using Versant Corporation's OpenAccess JDO implementation.

Results themselves, neither individually nor collectively, do not concern JEM's agents, except that they represent successfully completed chores associated with a particular crop, requiring persistent storage. The Farmer monitors crops, and determines the crops "ripeness" by comparing the completed chore count to the total chore count. Once the Farmer designates the crop as complete, other software can compile and distill completed crops outside of JEM.

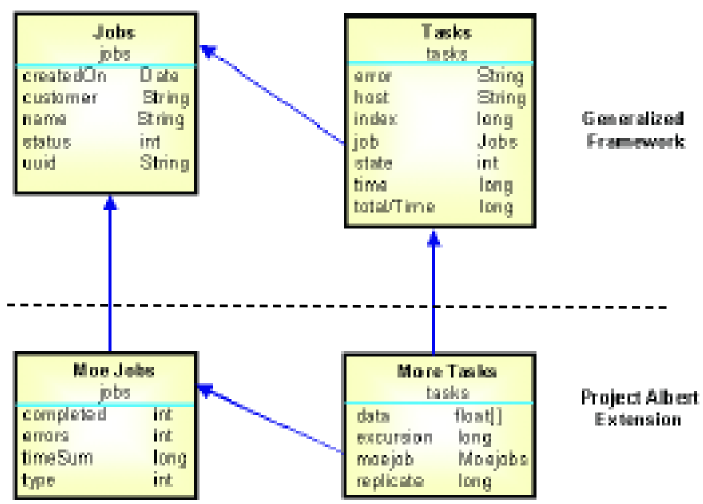

Figure 3. Persistent Repository.

Conclusion: JEM's agents provide a distributed enterprise grid computing service derived from the analogy of a farm. Every farm crop contains one or more chores performed by Worker agents. A Farmer agent distributes chores to Worker agents. A Manager agent distributes a crop's seed to the Farmer agent, distributes farming tools to both the Farmer agent and the Worker agents, and monitors the state of all the farm's attributes.

For reasons of robustness, reliability, scalability, and performance, JEM agents are not coupled, but instead communicate asynchronously across the network using either the JINI JavaSpace API or the JDO API, depending upon the nature of the message itself.

Significance: In contrast to "boutique" grid computing solutions, like SETI@home, *et. al.*, JEM provides a generalized solution to grid computing problems. The enterprise can use JEM's API to arbitrarily extend the computational grid to encompass any embarrassingly parallel computing requirement. Further, JEM's solution allows dynamic delivery of executable objects to its agents as the agents need them, allowing radically different types of jobs to be processed simultaneously. All of the required server software can run on the same modestly configured, low cost system without incurring performance liabilities. The lightweight Worker agent can be installed on every other enterprise CPU, creating a grid that can, when all Workers are activated, expand to the entire enterprise during times that the CPUs would otherwise be idle (overnight, weekends, and holidays). During business hours Workers may be deactivated when the CPUs are needed for routine desktop work.

JEM differs from other job schedulers used by MHPCC in JEM's use of a persistent repository to store information. This provides the ability to track CPU use over time using commercial off-the-shelf reporting software. It is easy to see the added value of JEM in the ability to report use by customer, use by CPU time accumulated by system, *et c.*

JEM's persistent repository also provides a means of examining system performance over time. Through graphing utilities developed by Robert Swanson at MHPCC and using performance metrics collected by JEM, it is easy to discover performance aberrations of individual systems compared to their peers. In one recent incident, three systems failed within two weeks of discovering aberrant performance trends during farming. The ability to discover potential system failure before it actually occurs offers a significant value over the other schedulers currently used by MHPCC.

I recently extended JEM's API to provide an XML based web service layer, exposing the Manager agent's operations for use via the HTTP protocol with SSL encryption. The Service Oriented Architecture (SOA) demonstrated its utility during the Project Albert International Workshop held during May 2005, in Stockholm, Sweden. Conference attendees sent 42 jobs from Stockholm to 44 CPUs on Maui (twelve hours and half the globe away), accumulating 38 CPU-days of remotely accessed and controlled simulations.

Planning the expansion of the SOA to encompass the Worker agent is complete. Upon completion of the implementation, the enterprise will be able to create an uncoupled, secure computing grid that includes federations of CPUs from sites scattered across the enterprise: state-wide, nation-wide, or world-wide.

Author and Contact: Michael S. Coulman
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Resources: Pentium 4 Mobile Sony Vaio Notebook at MHPCC
Sponsorship: Marine Corps Warfighting Laboratory (MCWL)

Web-Based Secure HPC Access

Brian Kruse, Brian Banks, Aaron Culliney, Maria Murphy, Mark Radleigh

Two important High Performance Computing Modernization Program (HPCMP) user requirements are addressed by this project: 1. A robust Kerberos/SecurID authentication mechanism for web servers and 2. A general-purpose remote XWindows desktop service that provides good performance over a Wide Area Network (WAN).

Research Objectives: MHPCC supports the STEPNet project, which requires Kerberos/SecurID authentication but does not use standard remote-access protocols (i.e., ktelnet, ssh, krsh). Specifically, users require XWindows sessions started on HPCMP systems to be remotely and securely displayed at their desktop systems. Although this can be accomplished via SSH,[1] sluggish performance and lack of resiliency that X-over-SSH sessions currently provide are not acceptable for STEPNet users. To address this problem, MHPCC software engineers developed two software packages: a web-enabled HPCMP Kerberos/SecurID authentication mechanism and a modified version of VNC.[2] The resulting application is a Kerberized VNC session displayed at the user's desktop via an SSL-enabled Java web interface. This solution has achieved a 10x performance improvement over current SSH X11 tunneling methods.

Kerberos/SecurID Authentication for Web Servers

Problem statement: Current access control methods for web servers do not incorporate NRL Kerberos/SecurID authentication. In order to secure non-public HPCMP websites, administrators currently use separate methods. We developed a module for the Apache web server that utilizes NRL Kerberos/SecurID for basic authentication.

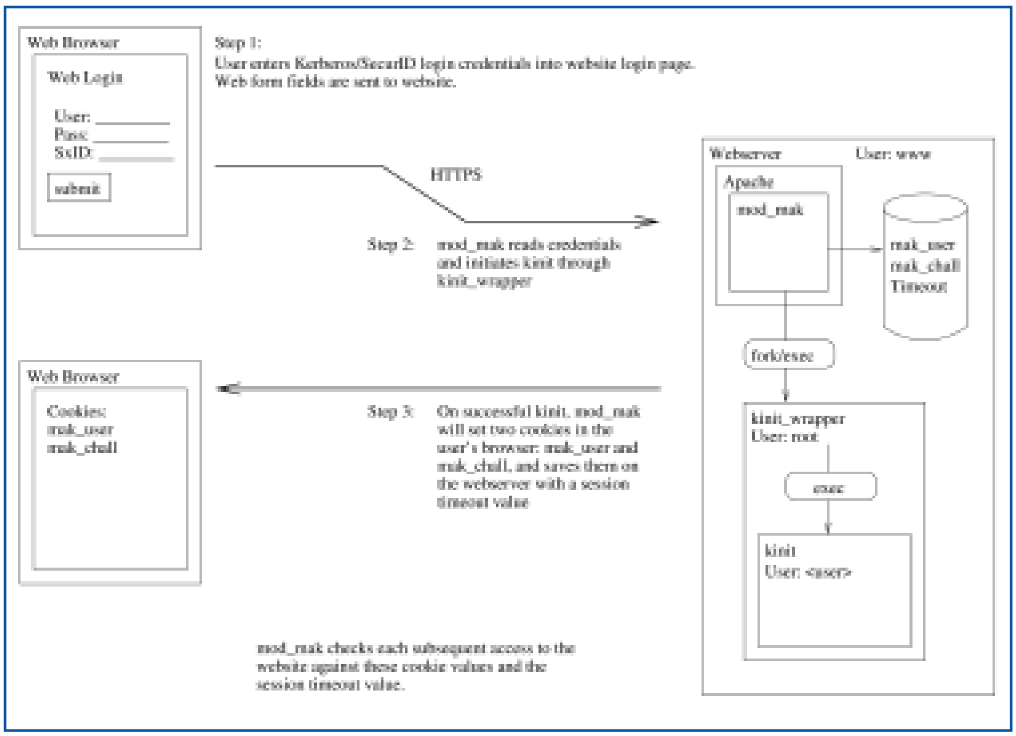

Figure 1. Mod_Mak Data Flow.

Reference Notes:

[1]The Secure Shell (SSH) is a protocol for creating a secure connection between two systems. In the SSH protocol, the client machine initiates a connection with a server machine. The following safeguards are provided by SSH:

- *After an initial connection, the client verifies that it is connecting to the same server during subsequent sessions.*
- *The client transmits its authentication information to the server, such as a username and password, in an encrypted format.*
- *All data sent and received during the connection are transferred using strong 128-bit encryption, making it extremely difficult to decrypt and read.*
- *The client has the ability to use X11 applications launched from the shell prompt. This technique, called X11 forwarding, provides a secure means to use graphical applications over a network. [We found this method to perform sluggishly in our environment.]*

[2]The Virtual Network Computing (VNC) protocol is a platform independent protocol that allows users to interact with the desktop of a remote host. The VNC protocol is a simple protocol for remote access to graphical user interfaces. It is based on the concept of a Remote Framebuffer (RFB). The VNC protocol simply allows a server to update the framebuffer displayed on a viewer. Because it works at the framebuffer level, it is potentially applicable to all operating systems, windowing systems, and applications. The protocol will operate over any reliable transport such as TCP/IP. This is a truly "thin-client" protocol: it has been designed to make very few requirements of the viewer.

Mod_Mak: This Apache Web Server module implements the HPCMP NRL Kerberos/SecurID authentication mechanism over SSL.[3] The web server utilizes the PIPE cache to perform the initial authentication and does not retain the Kerberos credentials. Instead, a separate 128-bit MD5 challenge is generated and associated with a web client, along with a web session timeout value that follows the normal procedures for the Kerberos 'authtime' attribute (i.e., 1 hour or 10 hours - depending on what is set in /etc/krb5.conf).

The mod_mak Apache web server module supports server-side authentication and authorization of web access via Kerberos/SecurID. The mod_mak module prevents access to a website until a user has been properly identified as an HPCMP-authorized user. When used in conjunction with Secure HTTP (HTTPS) via TLS,[4] mod_mak provides a comprehensive solution to securing HPCMP web resources.

A General-Purpose Remote XWindows Desktop Service

Problem statement: Running XWindows applications over a WAN is problematic because the X protocol is not sufficiently secure, efficient, and robust over a network. In other words, running X over the Internet is potentially dangerous, slow, and unreliable. Currently, there are only partial solutions to these problems, like tunneling the X protocol over SSH, which can yield sluggish performance. To address this problem, we developed a comprehensive, robust, and fast solution for running XWindows applications over a WAN.

VNC/KRB5: This package allows secure VNC connections to HPC resources over a WAN. Specifically, in conjunction with the HPCMP Kerberos/SecurID authentication standards, this package can enable a remote XWindows desktop display to the user's terminal utilizing Secure Sockets Layer/Transport Layer Security (SSL/TLS) and Java technologies.

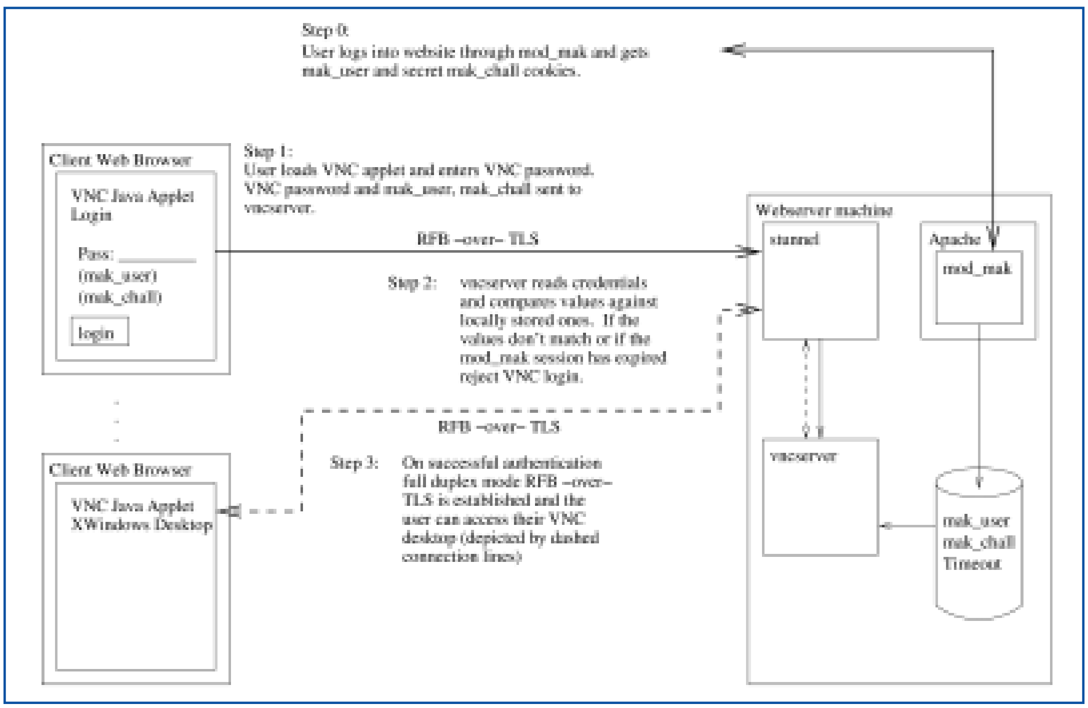

Figure 2. VNC/KRB5 Data Flow.

Reference Notes:

[3]Secure Sockets Layer (SSL) is a protocol developed by Netscape for transmitting private documents via the Internet. SSL works by using a private key to encrypt data that is transferred over the SSL connection. Both Netscape Navigator and MS Internet Explorer support SSL, and many web sites use SSL to obtain confidential user information, such as credit card numbers. By convention, URLs that require an SSL connection start with https: instead of http:.

[4]Transport Layer Security (TLS) is the successor to SSL. TSL is based on SSL version 3.0; however, TLS and SSL are not interoperable. The TLS protocol does contain a mechanism that allows TLS implementation to back down to SSL 3.0. TLS is composed of two layers: the TLS Record Protocol and the TLS Handshake Protocol. The TLS Record Protocol provides connection security with some encryption method such as the Data Encryption Standard (DES). The TLS Record Protocol can be used without encryption. The TLS Handshake Protocol allows the server and client to authenticate each other and to negotiate an encryption algorithm and cryptographic keys before data is exchanged.

The STEPNet project users require XWindows sessions started on HPC systems to be remotely and securely displayed to their desktop systems. We observed that using a traditional approach (i.e., tunneling X over SSH) does not yield acceptable performance. Whenever the screen needs to be updated (mouse move; window move, resize, open; etc.), the response time is very slow, and, at times, one can almost observe XWindows widgets rendered one at a time in slow motion. Our VNC solution works much faster and provides a swifter and more responsive user interface.

The VNC software provides the user with a remote XWindows desktop accessed through a Java applet. The user can manipulate remote XWindows programs from their local Microsoft Windows desktop without requiring installation of an XServer-for-MS Windows software like Exceed. The VNC Remote FrameBuffer (RFB) network protocol is very efficient and reliable, making it especially useful when WAN access is required. Even when tunneled over SSL/TLS, the combined RFB+TLS protocol is faster than other equivalent access methods, like XWindows tunneled over SSH or TLS. In addition, the VNC RFB protocol is ideal for dealing with inevitable network outages. After a prolonged network outage, the user only has to relogin to the remote VNC service to get right back to all his XWindows programs (remaining right where he left them on his desktop when access was lost). This compares quite favorably to having a raw or tunneled XWindows protocol fail completely during network outages, sometimes leading to data loss, and always requiring restarting every XWindows program. The access to VNC is secured by tunneling the raw RFB protocol over TLS, and authentication requires successful login via the mod_mak Apache web server module, based on NRL Kerberos/SecurID.

The net result is that this solution provides remote XWindows users the same capabilities and access on HPC assets as local users, and it does it in an efficient manner and with good user interface performance.

Authentication Sequence

The users view their VNC desktop through a Java applet running in a web browser (e.g., Microsoft Internet Explorer, or Netscape Mozilla). The authentication mechanism to set up this secure channel is a multi-step process:

1. The user logs in a secure website that uses the mod_mak Apache web server authentication module. The mod_mak module delegates authentication to NRL Kerberos/SecurID.
2. After successful authentication to the website, the mod_mak module stores secret access credentials in two temporary volatile cookies in the user's browser (mak_user, mak_chall). At this point the user can load the VNC Java applet on a separate page of the website.
3. The user must type in a secondary VNC authentication password to the applet to initiate the connection to the vncserver. The applet sends this password and the recently established mod_mak credentials over a secure TLS connection.
4. The vncserver checks the VNC password and the mod_mak credentials sent to it by the applet against local copies of them stored securely in a file (readable only by root and the user in question). If all authentication tokens match and the mod_mak session has not expired, the RFB session is established and the user receives access to the XWindows desktop. Otherwise, the connection is rejected.
5. This process ties the VNC session to a successful login via NRL Kerberos/SecurID through the mod_mak Apache web server module. When the user quits the browser, the mod_mak session cookies are deleted, and the entire authentication process must be repeated to reestablish the connection to the desktop.
6. If the user does not quit the browser, the VNC session will expire when the mod_mak session times out. At this point, the VNC network connection is dropped and can only be reestablished by reauthenticating through mod_mak. This session expiration value is ultimately determined by the Kerberos/SecurID installation and the settings in the /etc/krb5.conf file.

Acknowledgements

The software solution described in this briefing was developed to support the Simulation Test and Evaluation Process Network (STEPNet) project. STEPNet is funded by the HPCMP under their Portfolio program [Topic Area: Collaborative Simulation and Test (CST), Portfolio Area: Portfolio Integration Framework (PIF)]. The STEPNet project is a joint effort between NAVAIR, Boeing's Rocketdyne Propulsion and Power Division, and MHPCC.

Author and Contact: Brian Kruse
Authors: Brian Banks, Aaron Culliney, Brian Kruse, Maria Murphy, Mark Radleigh
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Resources: *Koa* and *Huinalu* Linux Clusters at MHPCC
Sponsorship: DoD High Performance Computing Modernization Program (HPCMP)

NASA Helios Mishap Investigation

Kevin Roe, Duane Stevens, John Porter

The NASA solar powered research aircraft, Helios, crashed approximately 10 NM west of the Pacific Missile Range Facility, Kauai, Hawaii on June 26th, 2003. The mishap cause is unknown but believed to be due to turbulent conditions. Important to this investigation was a thorough understanding of the atmospheric conditions during flight. Unfortunately, no atmospheric observations were available in the vicinity of the mishap. However, the flow of winds around nearby mountains east of PMRF can cause "shear lines", marking an enhanced vertical component of airflow vorticity, between trade-wind flow diverted around Kauai in the vicinity of PMRF, creating wake-type secondary circulations to the lee of Kauai in the vicinity of PMRF land areas where some atmospheric observations are available. Consequently numerical simulations using very high-resolution atmospheric computer models are required to characterize offshore conditions.

Research Objectives: Provide the NASA Helios Mishap Investigation Board information about overall flow patterns in the vicinity of Kauai and airflow turbulence, induced by flow over and around the island. The turbulence in these regions has never been characterized quantitatively as to its structure and intensity. Even the nature and detailed structure of the frequent nearby shear lines have never been well documented, nor is information available regarding the sources of gustiness, such as mountain waves or elevated shear layers that might be relevant. Combining computer simulations of conditions on the mishap day with actual observations of the shear lines from an instrumented aircraft, even though from other days, will greatly increase our understanding of the possible atmospheric conditions that may have contributed to the mishap.

Figure 1. NASA Helios prototype flying over Kauai, HI.

Results: Modeling studies with the MM5 model were carried down to the 1 km resolution and showed the location of the calm zone in the lee of Kauai. The model also showed that the synoptic conditions were fairly typical trade winds, with north easterlies on the south turning to east north east to the north side of the island chain. Both the MODIS satellite image and the MM5 1 km model showed a calm region in the lee of the island, surrounded by stronger trade winds. The calm zone and the north and south shear lines were observed in the MODIS image to be oriented straight westward from Kauai. The model calm zone was oriented slightly to the north of west. Model results indicated a calm zone bounded by shear lines on the north and south edges. These shear lines were vertically oriented without horizontal tilt at the longitude of the incident.

Based on numerical model results and instrumented aircraft findings, there are regions of strong shear around the island. This shear line, coupled with turbulent fields in the lee of Kauai, is a potential reason for the crash. Further study needs to be conducted including: 1) additional flights to obtain a better picture of semi random turbulent features, 2) ensemble modeling methods may provide a probabilistic estimate of turbulent and ambient fields, 3) real time measurements can aid in risk avoidance and better estimate day to day estimates of turbulent conditions.

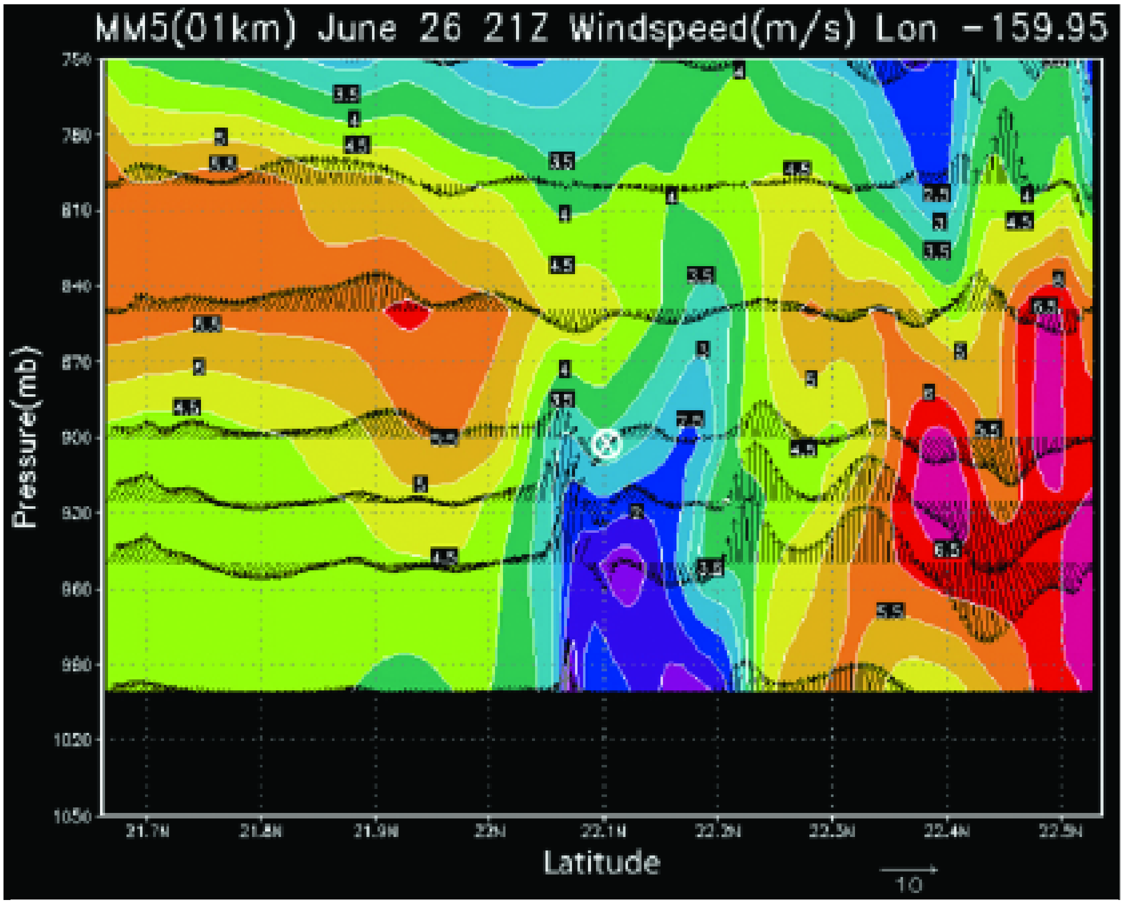

Figure 2. Windshear over Kauai, HI.

References:

1) J. Michalakes, "The Same-Source Parallel MM5," *Journal of Scientific Programming*, 8, 2000: 5-12.
2) J. Michalakes, J. Dudhia, D. Gill, T. Henderson, J. Klemp, W. Skamarock, and W. Wang, 2004: "The Weather Research and Forecast Model: Software Architecture and Performance,"to appear in proceedings of the 11th ECMWF Workshop on the Use of High Performance Computing In Meteorology, 25-29 October 2004, Reading U.K. Ed. George Mozdzynski.
3) J. Michalakes, J., R. Loft, A. Bourgeois (2001): "Performance-Portability and the Weather Research and Forecast Model" in on-line proceedings of the HPC Asia 2001 Conference, Gold Coast, Queensland, Australia, September 24-28, 2001.
4) K. P. Roe and D. Stevens, "High Resolution Weather Modeling in the State of Hawaii," The 11th PSU/NCAR Mesoscale Model User's Workshop, Boulder, CO, 2001.

Author and Contact: Kevin P. Roe
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Author: Duane Stevens
Organization: Department of Meteorology, University of Hawaii at Manoa, 2525 Correa Road, HIG 350, Honolulu, HI, 96822
Author: John Porter
Organization: Hawaii Institute of Geophysics and Planetology, University of Hawaii at Manoa, 2525 Correa Road, HIG 350, Honolulu, HI, 96822

Resources: IBM SP3/SP4 at MHPCC
Sponsorship: NASA/Kennedy Space Center

Intelligence Fusion and Semantic Fusion Publishing

Greg Seaton, Marc Lefebvre, David Norman

The Maui High Performance Computing Center (MHPCC) designed and developed an intelligence fusion system for the Office of Naval Intelligence (ONI) to solve the problem of accessing and utilizing multiple intelligence sources accessible through different systems and interfaces (i.e., "stovepipes"). The intelligence fusion system combines, or fuses, intelligence from multiple, disparate sources into a single, coherent interface. The intelligence fusion process implements a universal ontological framework across the distinct sources, leveraging object-oriented technologies to provide fused intelligence to human analysts and artificial intelligence (AI) agents with the goal of generating high-value, actionable intelligence for the intelligence consumer. In addition to decision support, the intelligence fusion architecture allows subject matter experts (SME) to capture their valuable subject matter expertise to be leveraged by other analysts and AI agents. Semantic fusion publishing allows access to fused intelligence by publishing fused entity intelligence to semantic artifacts, reducing duplicate and intensive fusion processes and achieving better distribution of fused intelligence to analysts.

Background: Intelligence fusion has emerged to include any combination, or fusing, of multiple intelligence sources from a single request into a single result. This type of generalized, analyst-focused intelligence fusion includes retrieval of all relevant information about a particular entity from all accessible intelligence sources.

An example of intelligence fusion would be returning all of the intelligence from accessible sources about a particular person of interest possibly including personal information, associations, affiliations, criminal record, financial records, medical history, and/or a timeline of movements.

Intelligence Fusion vs. Data Warehousing

Intelligence fusion is quite different from data warehousing. The data warehouse approach forces all data accessible to the end-user to be first processed through an extraction-transform-load (ETL) procedure and stored in a centralized and monolithic database. Intelligence fusion only accesses and retrieves the intelligence necessary to satisfy a request and does not require archiving or storing the resulting fused intelligence. The different data query and intelligence request approaches are illustrated in Figure 1a and Figure 1b.

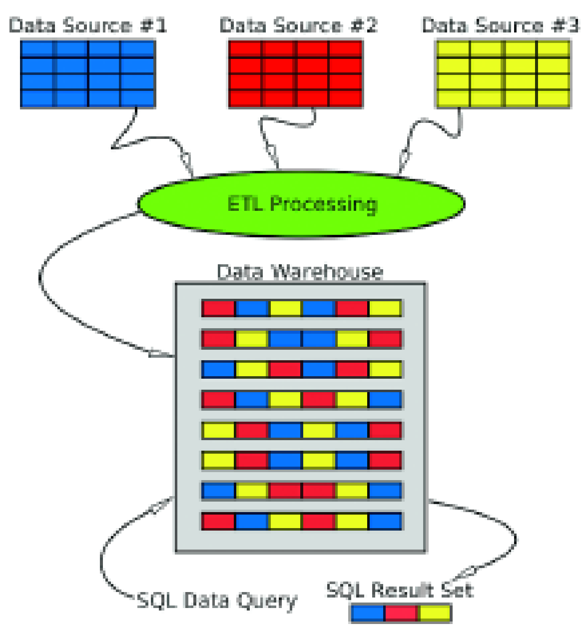

Figure 1a. Data Queries in Data Warehouse Systems.

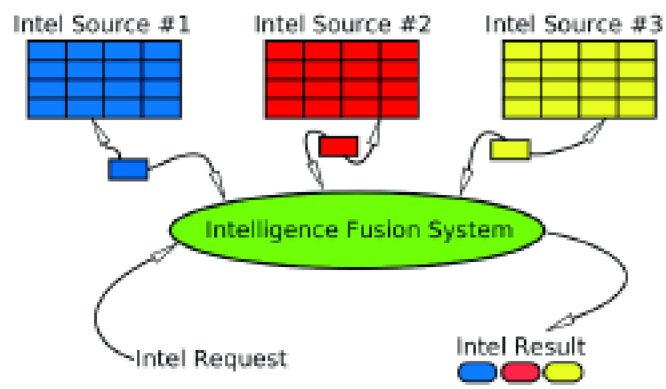

Figure 1b. Intelligence Requests in Intelligence Fusion Systems.

There are many advantages to using a intelligence fusion design over a data warehousing approach.

In a data warehousing effort, all of the data from the data sources must be extracted from the data sources, transformed into structures compliant with the data warehouse structures, and loaded into the data warehouse. The resulting data warehouse requires storage space that will grow linearly throughout the life cycle of the warehousing project, requiring increasingly more administration and storage resources along with the complicated issue of data ownership and control.

Core Intelligence Fusion System Components

The intelligence fusion system consists of four (4) primary components in a services-oriented architecture (SOA):

- Fusion Ontology Workspace
- Subject Matter Expert Workspace
- Ontology Service
- Intelligence Access Service

Fusion Ontology Workspace
The fusion ontology workspace GUI client application allows the fusion analysts and administrators to map element instances across disparate intelligence sources, define and extend element ontologies, and establish and administer security across the intelligence fusion system.

Subject Matter Expert Workspace
The subject matter expert (SME) workspace GUI client application allows the SME to capture her expertise as formal rules. These rules may then be used by AI agents as well as by other analysts.

Ontology Service
The ontology service manages all ontological operations, including creation, modification, usage of the system ontologies, expert rules, and ontological mappings.

Intelligence Access Service
The intelligence access service manages all connections, access control, and intelligence retrieval for the intelligence sources in the intelligence fusion system.

These four (4) components (Figure 2) comprise the core architecture and functionality of the MHPCC intelligence fusion system with varied analyst workspace service-request clients.

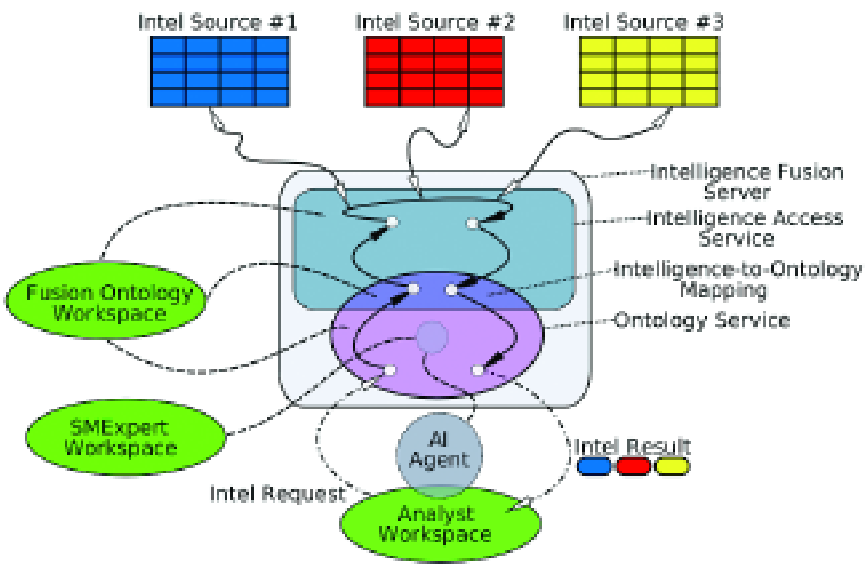

Figure 2. Core Intelligence Fusion System Components.

AI Agents

The MHPCC intelligence fusion system is architected to leverage multi-tiered AI agents across the collected intelligence sources to provide true 24/7 utilization of system resources. Expert AI agents, utilizing formalized subject matter expertise captured from SME analysts, may spider the ontological topography to identify and flag entities of interest for further automated analysis and/or consideration by human analysts.

Since the intelligence may be incomplete for any given entity and the amount of time to process the intelligence is not infinite, the intelligence fusion system uses the concepts of limited rationality (acting on incomplete intelligence) and satisficing (making decisions in less time than calculating complete optimal solutions) to provide a pragmatic and practical implementation of AI agent technology.

Semantic Fusion Publishing

Semantic fusion publishing allows access to fused intelligence by publishing fused entity intelligence to semantic artifacts, reducing duplicate and intensive fusion processes and achieving better distribution of fused intelligence to analysts.

The benefits of XML Topic Mapping and RDF/S semantic publishing (Figure 3) are many-fold, including:

- XHTML documents (XML transformed by XSLT to XHTML) presentation to the users using standard web browsers
- XML documents archive for possible further processing and searching (XQuery, etc.)
- Topic maps for simplified searching and navigation
- RDF/S for supporting AI agents in a distributed environment

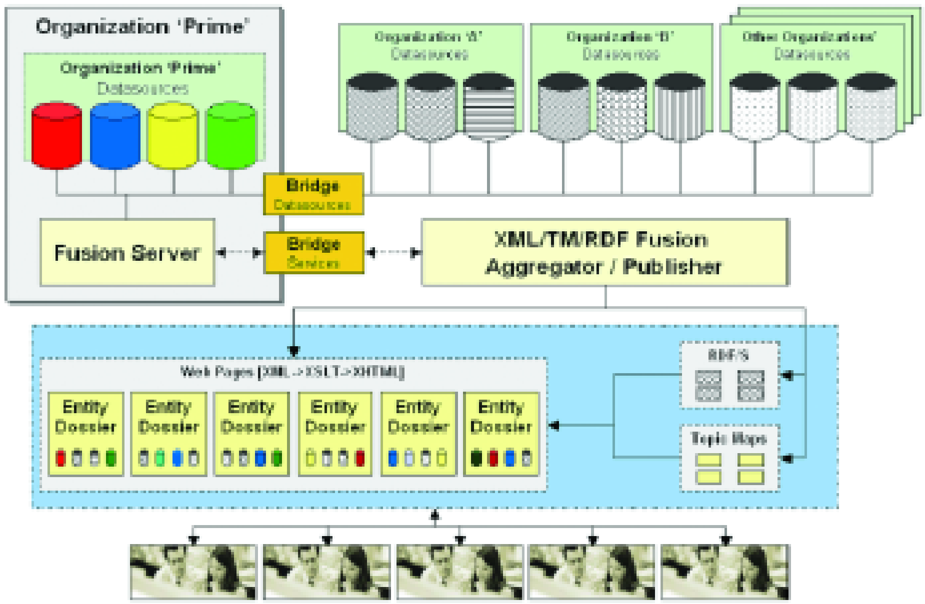

Figure 3. Semantic fusion publishing with topic maps and RDF/S.

Summary: By leveraging an object-ontological hybrid approach, the intelligence fusion system provides the benefits of presenting the underlying intelligence topography as cogent and coherent objects to the end-user and leverages the power and flexibility of an ontological, knowledge-oriented framework.

In addition, the MHPCC intelligence fusion architecture allows for the implementation of HPC/GC approaches to provide a high-availability and extensible environment, optimizing computational resource usage and maximizing the amount of analytical processing on the target intelligence domain.

The MHPCC intelligence fusion system provides an interactive and productive solution for analysts and agencies requiring a single, coherent interface to disparate, heterogeneous intelligence sources.

The goal of the MHPCC intelligence fusion system is to allow analysts to spend more of their valuable time analyzing high-interest entities and less time fighting the interface, ultimately resulting in more high-value, actionable intelligence for the intelligence consumer.

Author and Contact: Greg Seaton
Authors: Marc Lefebvre and David Norman
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753

Resources: IBM SP Power4 Systems and *Huinalu* Linux Cluster at MHPCC

Sponsorship: Office of Naval Intelligence (ONI)

The Unmanned Systems Test Bed (USTB)

Brian Kruse, Brian Banks, Jeff Beck, Lawrence Brem, Matthew Burnham, Richard Cook, Michael Coulman, Jonathan Dann, Scott Hofmann,Thomas Meyer, Barry Pang, Michael Smith, Shannon Wigent

The primary objective of the Unmanned Systems Test Bed (USTB) program is to develop new training, test, and evaluation concepts and capabilities for unmanned systems through a series of demonstrations. The Defense Test Resource Management Center (DTRMC) has teamed up with the Defense Information Services Agency (DISA) to examine how to augment DOD Test and Training Range resources with new capabilities in support of unmanned systems. DTRMC is interested in finding effective uses for advanced visualization software and high-performance computing assets to support the needs of unmanned systems. Related goals include leveraging the Test and Training Enabling Architecture (TENA) for cross-range and cross-facility data collaboration, and running TENA in a parallel processing environment. Accordingly, the initial stages of the USTB program focuses on applying advanced visualization and high-performance computing technologies to support the needs of unmanned systems users.

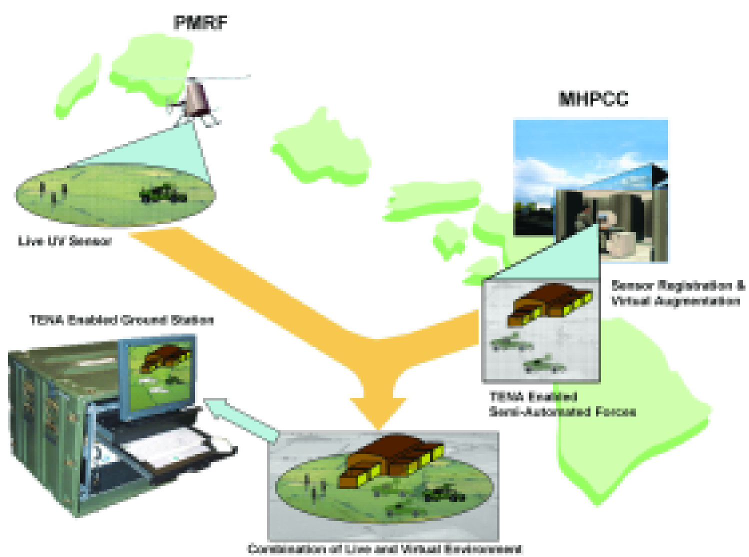

Figure 1. USTB General Concept.

Research Objectives: Unmanned vehicles play an increasingly important role in the war fighting systems of the U.S. military. Although the trend began with intelligence, surveillance, and reconnaissance (ISR) assets (e.g., Predator and Global Hawk), it now includes combat vehicles (e.g., Armed Predator and Air Force and Navy Unmanned Combat Air Vehicle) and utility platforms (Quick Delivery ACTD). The evolving network-centric warfare concepts of all services rely increasingly on unmanned platforms. The associated command, control, communications, computers, intelligence, surveillance, and reconnaissance (C4ISR) systems require interoperability standards that include unmanned vehicle mission planning, control, and data dissemination. Recent DOD R&D efforts reflect emphasis toward all types of unmanned systems: air, ground, sea surface, and underwater platforms. Transformational programs such as the Army's Future Combat System and Navy's SSGN point to future capabilities dependent on unmanned systems. Consequently, future operations may use unmanned vehicles at all echelons of command, requiring interoperability and integration of mixed unmanned vehicle types. It is increasingly important for the individual services and the Joint Commands to develop concepts and facilities for training exercises, CONOPS development, and system T&E, which include single and combined unmanned vehicle operations. Though distributed exercises involving computer-generated virtual environments and simulation accomplishes much of this objective, the exercises must also include live components on a test range that provides for air, ground, sea surface, and underwater unmanned vehicle operations.

The USTB Concept:

The key concepts behind this test bed are:

1. Capture three-dimensional terrain data of the test range.
2. Use the terrain data to render a realistic view of the unmanned system's environment. Simulate a virtual flythrough in real time, synchronized with the movement of the unmanned system.
3. Integrate any sensor data (e.g., video) from the unmanned system with the synthetic scenery. Project the live video like a spotlight on the rendered synthetic scenery, providing a window to the real world displaying any movement, change, or action in the observed area.
4. Integrate simulated entities (buildings, troops, vehicles, etc.) with the real environment to support various test, evaluation, and training scenarios.
5. Integrate these systems with standard battle space management software.

Figure 1 illustrates the USTB concept showing how the rendered synthetic scenery, live video stream captured by the unmanned system sensor, and real and simulated battle space entities combine to create a unified battle space environment.

The USTB System Components: Figure 2 shows a simplified view of the USTB system architecture supporting the general concept depicted in Figure 1. The unmanned system sends video and telemetry to the ground station at the Pacific Missile Range Facility (PMRF). PMRF then sends the video and telemetry to the Maui High Performance Computing Center (MHPCC) over the Defense Research Network (DREN) via an OC-12 (622 Mbits/sec) fiber link. MHPCC sends processed information to PMRF over the same data link to the ground station for display.

1. **Terrain data capture:** A LIDAR[1] system captures the three-dimensional terrain, yielding high-resolution data with accurate geospatial coordinate information.
2. **Terrain database management and scene generation software:** Scene generation software interactively renders the terrain data, video, and simulated entities from the appropriate viewpoint.
3. **Video projection and alignment software:** This component receives the video from the camera on the unmanned system and processes it for integration with the synthetic scenery.
4. **Battle space management software:** This software manages real and simulated entities. Its user interface is a map-like view that depicts the entities (troops, buildings, ground vehicles, aircraft, ships, etc.) as icons. The battle space management software display is separate from but adjacent to the scene generation software display (merged video and synthetic scenery).
5. **Test & Training Enabling Architecture (TENA) Middleware:** TENA middleware provides communication between all the software components. The TENA object model encapsulates all data shared by the system components.

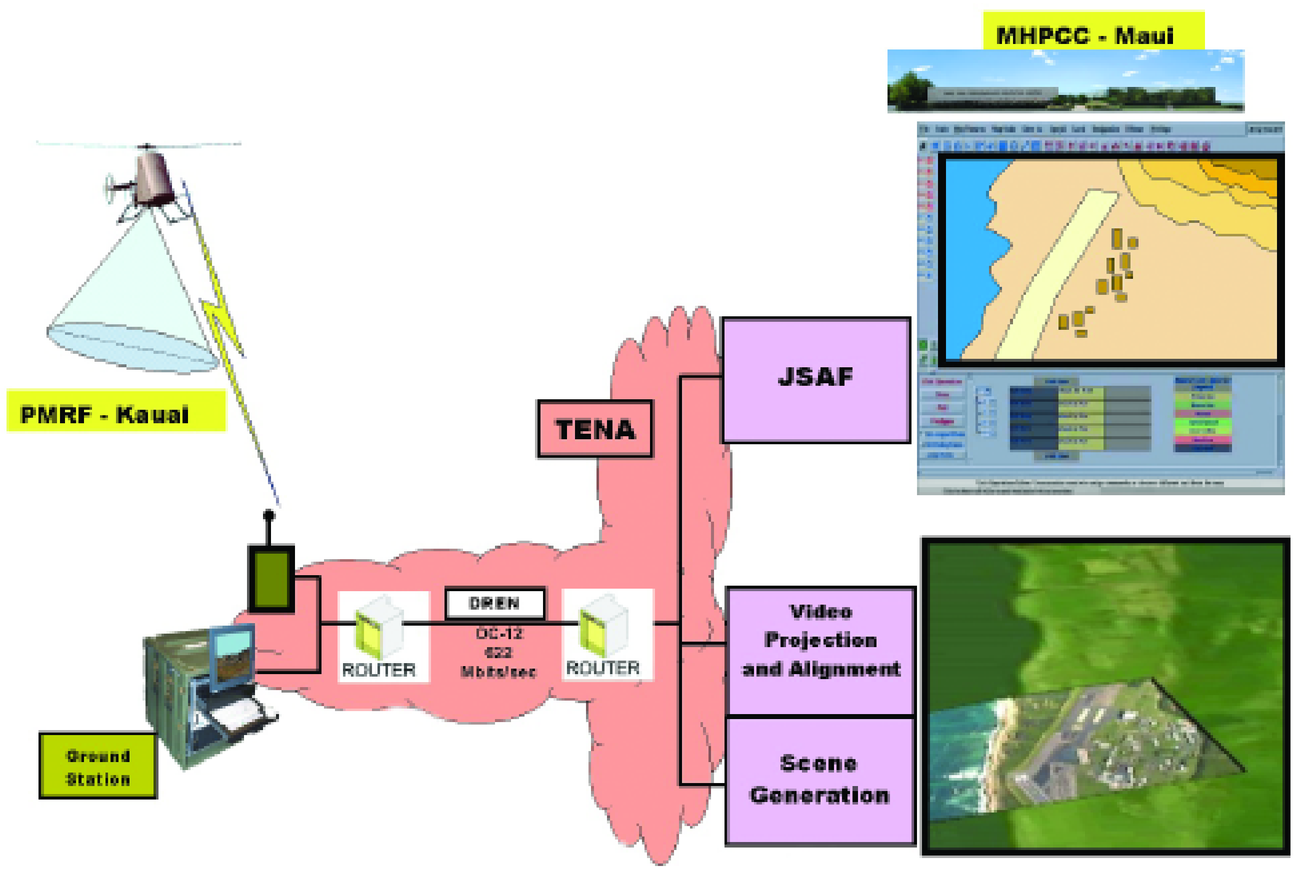

Figure 2. USTB System Architecture.

Reference Note:

[1]LIght Detection And Ranging uses the same principle as RADAR. The LIDAR instrument transmits light out to a target. Instruments analyze the transmitted light's reflection and scatter. The change in the properties of the light enables properties of the target to be determined. LIDAR uses the time for the light to travel out to the target and back to determine the target's range.

Methodology:

Video Projection: The goal is integration and rendering of the video with the synthetic scenery in real time, making the video appear as part of the synthetic scenery. The term video projection refers to using an algorithm to map the video onto a three-dimensional synthetic scene. The simplest projection may involve matching the four corners of the video frame to pixel locations on the synthetic scenery by using telemetry, then texture mapping the video frame to this area.

Video and Unmanned System Telemetry: Video projection uses telemetry data from the unmanned system. This includes the sensor telemetry (azimuth angle, zoom factor, etc.) and the unmanned system telemetry (GPS location, yaw, pitch, etc.). In general, a three-dimensional object (in this case the video camera) has six degrees of freedom: three coordinates x, y, and z and some kind of rotation, for example, three angles alpha, theta and gamma, describing its position and orientation relative to a frame of reference. In our case, the sensor telemetry provides the relative position of the camera with respect to the unmanned system. The unmanned system telemetry provides the position of the unmanned system relative to the earth. This is sufficient to calculate the position and orientation of the camera relative to the earth.

Video Alignment: The terms video alignment, video image registration, and video registration are often used interchangeably by industry, and refer to detecting and matching (i.e., aligning or registering) the image information from a video with a previously captured image of the same scenery. This usually involves image processing to detect shapes, edges, and features on the video frame and the previously captured image, matching the features from both, and aligning them. In our case, the video equivalent is the live sensor data from the unmanned systems and the previously captured image equivalent is the LIDAR data.

This can produce the following benefits:
1. Image processing the video frame and the synthetic scenery and finding edges, features, or objects common to both permits better alignment of the video and the synthetic scenery.
2. Information from the synthetic background enables accurate calculation of the location of objects observed in the video because the positional accuracy of the LIDAR terrain data is significantly better than that of the GPS on the unmanned system. This information allows derivation of the absolute location of the stationary objects and features (buildings, roads, etc.) in the video and determination of the relative location of the transient and moving objects (troops, vehicles, etc.) in the video relative to the permanent and stationary objects.

The primary goal is to align and register video frames (in real time) with rendered views of the scene derived from the LIDAR terrain dataset. A secondary goal is obtaining accurate coordinate information for the permanent and transient objects, features, and entities observed in the video by using a precisely aligned LIDAR terrain dataset.

Video Synchronization: The projected video and the synthetic scenery may also require time synchronization. The telemetry travels asynchronously as data packets over a network. There are multiple components in the loop from the unmanned system to the high-performance computing resources and back to the ground station, each capable of delaying transmissions. The network, hardware, and software may each add to these delays. If not compensated for, these delays and the asynchronous nature of the network traffic could result in the merged video and synthetic scene being off by many seconds and by 100s or 1000s of frames. This may yield unacceptable mismatches and artifacts (e.g., the synthetic scenery and the video may depict significantly different viewing angles).

The main challenge is accurately correlating the viewpoint for the synthesized scene and the actual sensor camera. Therefore, to achieve accurate synchronization between the synthesized image sequence and the video stream, the telemetry information (GPS data, camera angle, zoom factor, etc.) and the live video steam both must be stamped by the same time clock.

Author and Contact: Brian Kruse
Authors: Brian Banks, Jeff Beck, Matthew Burnham, Richard Cook, Michael Coulman, Jonathan Dann, Thomas Meyer, Barry Pang
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Authors: Lawrence Brem, Scott Hofmann, Michael Smith, Shannon Wigent
Organization: Science Applications International Corporation (SAIC), 5400 Shawnee Road, Suite 201, Alexandria, VA, 22312
Sponsorship: Defense Test Resource Management Center (DTRMC)

Theater UnderSea Warfare (TUSW)

Robert Dant, Michael Berning, Carl Holmberg, David Solomon, Chelsea Nesbitt, Thomas Meyer

The Air Force Research Laboratory Maui High Performance Computing Center (AFRL/MHPCC) maintains a Linux cluster for dedicated Theater UnderSea Warfare (TUSW) use. This TUSW cluster provides high performance computing resources to TUSW users for computationally intensive undersea warfare (USW) simulations. The cluster is designed to be easily expandable and has secure connectivity (via SIPRNet) to Pearl Harbor (Commander Task Force-12, CTF-12), allowing 'remote' access to the AFRL/MHPCC computing resources on Maui for computationally intensive acoustic modeling. TUSW software, such as STAPLE (Scalable Tactical Acoustic Propagation Loss Engine), AAT (Asset Allocation Tool), and Grid Manager, has been integrated to demonstrate "reachback" high performance computing during fleet exercises. Ocean environmental data has been provided via the TEDServices database. TUSW resources have been successfully demonstrated during Silent Fury (SF), RIMPAC, and Undersea Dominance (UD) exercises in the Pacific Theater, as well as during several "Local" CTF-12 (Pearl Harbor, Hawaii) training activities (Figure 1).

Research Objectives: TUSW program objectives for Year 3 activities at AFRL/MHPCC involved: 1) integrating software models and tools into the TUSW Linux cluster system, 2) developing and testing a TUSW Manager/Scheduler (TMS) system that would allow efficient use of the TUSW Linux processors during modeling activities, 3) providing shore-based high performance computing using operational databases and executing computationally intensive acoustic models, and 4) testing and evaluating the TUSW system during Naval exercises (Figure 2).

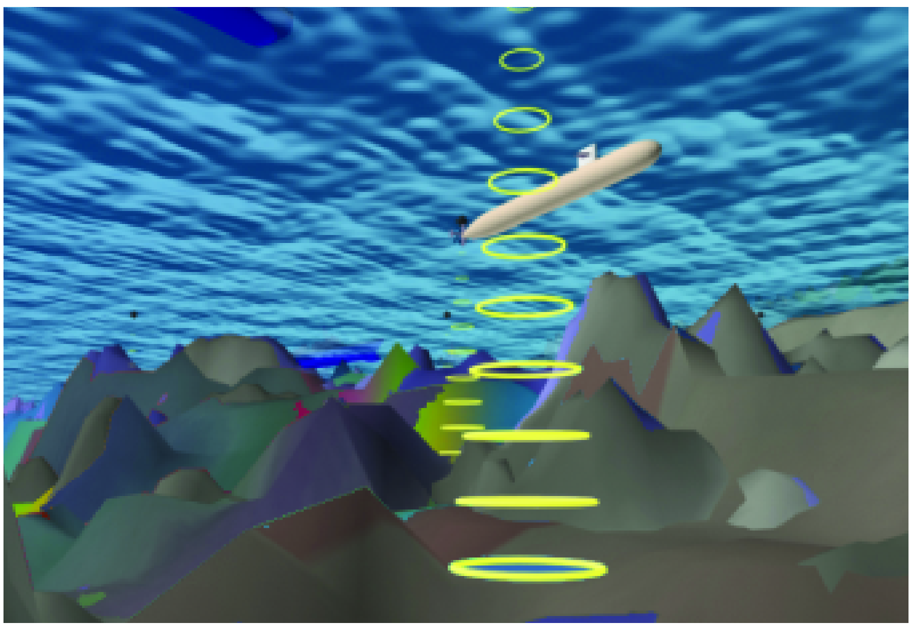

Figure 1. Display of simulated bathymetry, contacts, buoy patterns and time-lapse sonar data are part of the research to augment existing TUSW graphical software.

Methodology: Research efforts for the TUSW program (Year 3) at AFRL/MHPCC focused on: 1) supporting 'real world' Naval exercises (e.g., Silent Fury, RIMPAC, Undersea Dominance), 2) installing and testing upgrades to the acoustic modeling software on the TUSW cluster , and 3) providing secure connectivity (via SIPRNet) to CTF-12 (Pearl Harbor, Hawaii), CTF-74 (Yokosuka, Japan), and the Tactical Support Center (TSC) at Marine Corp Base Hawaii (MCBH, Oahu), 4) developing a TMS capability to distribute and manage TUSW applications on the TUSW cluster, and 5) providing upgrades, patches, and security enhancements to the TUSW cluster operating system (OS), to ensure an efficient, reliable, and secure operating environment for TUSW users.

AFRL/MHPCC successfully integrated several sets of enhanced software tools into the TUSW system, including an upgraded version of the Scalable Tactical Acoustic Propagation Loss Engine (STAPLE) to calculate sound transmission losses through the ocean, and a Tactical Environmental Database Services (TEDServices) capability that provided continual updates to STAPLE's ocean environmental data (Figure 3).

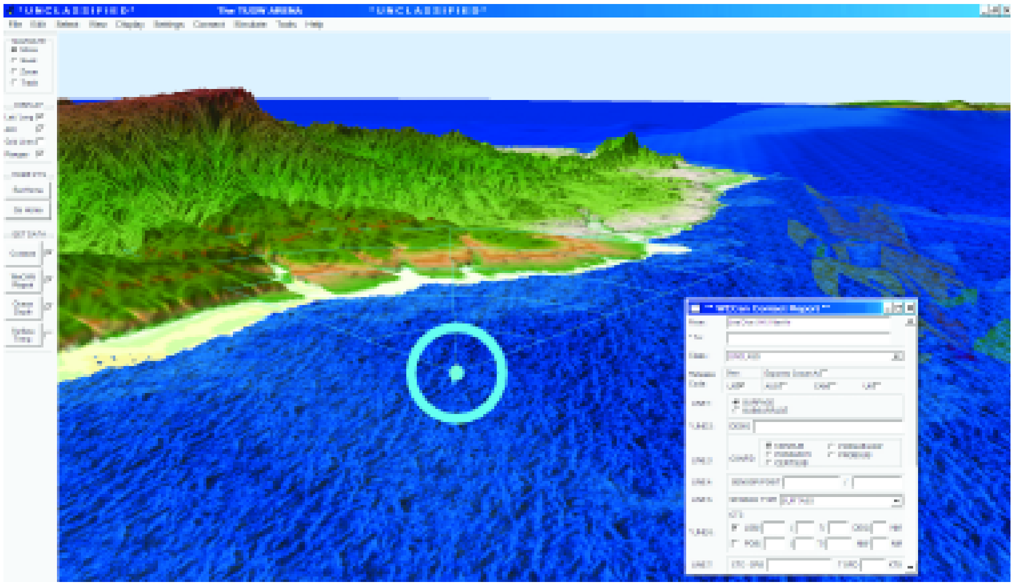

Figure 2. Ongoing TUSW research activities involved display of real world operational database information, such as the WeCAN (Web-Centric ASW Network) reports.

Results: In Year 3 of the TUSW program, AFRL/MHPCC expanded the TUSW cluster to a 48-node (96-processor) system for exclusive TUSW use, as well as provided a stand-alone server/workstation and substantial disk storage. Upgrades to the STAPLE transmission-loss modeling software were successfully installed, its performance characterized, and then integrated with network access middleware to permit computational 'reachback' by remote PC-based TUSW client workstations in the Pacific Theater.

AFRL/MHPCC also implemented a TUSW Manager/Scheduler (TMS) system that was responsible for managing remote user requests from CTF-12 (Pearl Harbor, HI) and CTF-74 (Yokosuka, JP). The TMS translated user requests into appropriate TUSW application requests, which were presented to the TMS for execution. The TMS, which was designed for high-volume, short-duration jobs demanded by the TUSW applications, then scheduled and orchestrated the execution of these acoustic modeling jobs on the TUSW cluster. The design of the TMS was ‘concurrent’ in nature, so as to maximize job throughput and cluster resource utilization. The TMS also integrated knowledge of application resource demands and database bathymetry synchronization, so as to provide a robust and data-centric computing environment.

In addition, AFRL/MHPCC implemented a secure network interface (via SIPRNet) between the TUSW Linux cluster and the 'remote' TUSW users at CTF-12 and CTF-74. AFRL/MHPCC also participated in TUSW code testing and evaluation, as well as developing proof-of-concept demonstrations that highlighted TUSW capabilities.

Future TUSW developmental efforts will include additional testing of the high performance computation 'reachback' concept during fleet exercises, enhancements to the TMS, performance benchmarking, and installation of additional enhanced/upgraded TUSW software models and tools.

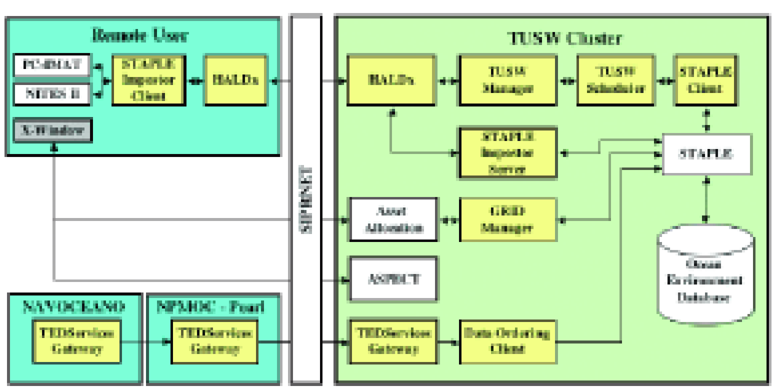

Figure 3. Acoustic models, databases, and clients are linked over a wide area network using networking middleware that brings high performance computing to deployed naval units throughout the Pacific Fleet.

AFRL/MHPCC Visualization efforts in Year 3 of the TUSW program included the development of a multi-platform interactive software application for viewing graphical representations of geo-referenced data within a 3-D geo-spatial environment (Figure 4). Built into an interactive viewer were tools for defining and zooming to a geographic Area of Interest (AOI), using TUSW contacts and targets having Naval Tactical Data System (NTDS) symbology and models. Terrain and bathymetric data were converted into geometries having multiple levels of detail and draped with geo-referenced aerial imagery. These databases represented both historical and concurrent measurements, such as bathymetry, bottom loss, sound velocity, ocean temperature, and potentially a multitude of other available measurements.

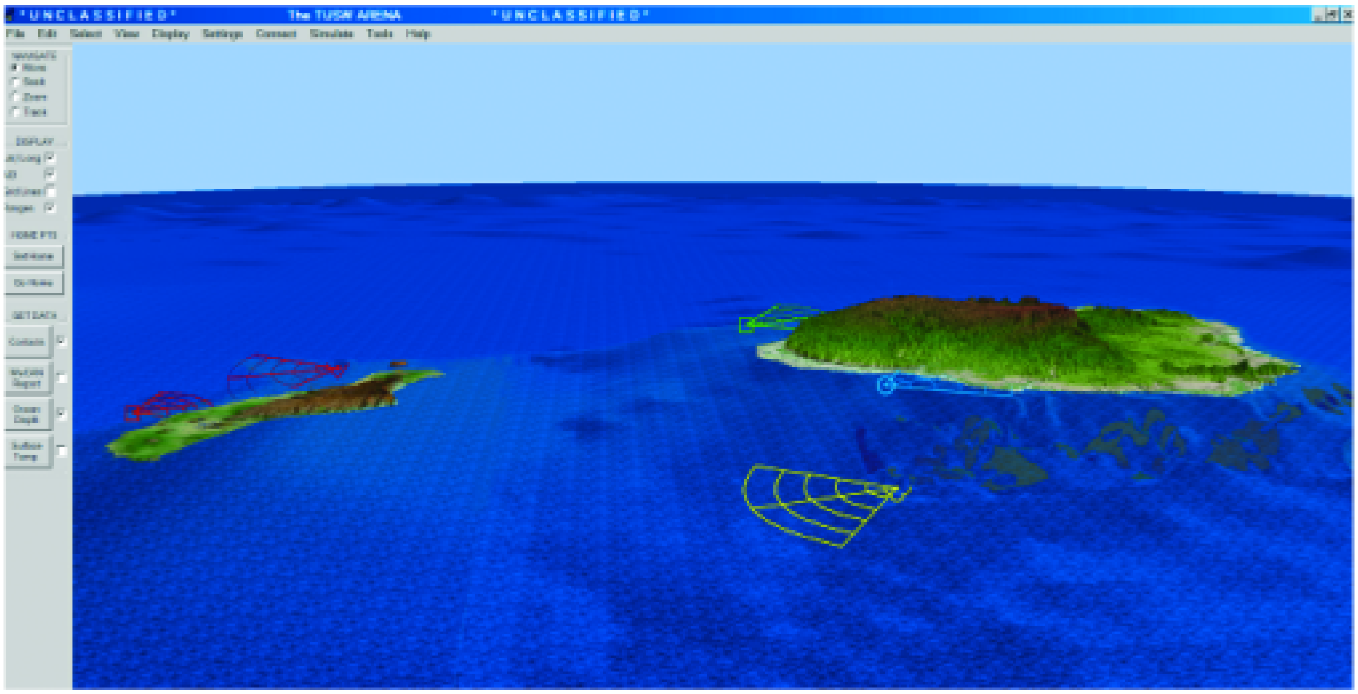

Figure 4. 3-D geo-spatial data environments, fused with interactive manipulation of contacts displayed with NTDS (Naval Tactical Data Systems) symbology, provide better visualizations of undersea warfare (USW) scenarios.

Future TUSW visualization activities at AFRL/MHPCC will include the live integration of USW data obtained on demand from active USW databases to enable multiple-asset tracking capabilities and near-real-time display of current oceanographic measurements for given global or local areas of interest.

Significance: The Office of Naval Research (ONR) has promoted research and development of innovations that will provide technology-based options for future U.S. Navy and Marine Corps capabilities. The goal for TUSW is to "Implement a 'leap ahead' command center for the future, for knowledge development, environmental analysis, and resource allocation for undersea warfare." These new technologies will be transitioned to the acquisition community for eventual integration into the Navy's operational USW fleet and the Marine Corp's warfighting inventory.

Author and Contact: Robert Dant
Authors: Michael Berning, Carl Holmberg, David Solomon, Chelsea Nesbitt, Thomas Meyer
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Resources: IBM Linux Cluster (96 Intel Processors), SGI Onyx 3900 (12 Processors) with Infinite Reality3 Graphics, and DELL Dual Xeon (2 Processors) with nVidia QuadroFX 200G Graphics at MHPCC
Sponsorship: Office of Naval Research (ONR)

A Theoretical Study on the Substitution on 2- and 3-Dimensional Aromatic Systems

Eluvathingal D. Jemmis, Biswarup Pathak, Anakuthil Anoop

DFT studies (B3LYP/6-31G*) on mono- and dichloroderivatives of benzene, naphthalene, $B_{12}H_{12}^{2-}$, four atom sharing condensed systems $B_{20}H_{16}$, and mono-carborane isomers of $B_{20}H_{16}$, are used to compare the variation of relative stability and aromaticity between condensed aromatics.

Research Objectives: The attempt is to study the substituent effect on condensed 2- and 3-D aromatic systems and make a comparison between themselves. Benzene and $B_{12}H_{12}^{2-}$ are important prototypes of two- and three-dimensional aromatic compounds in the carbon and boron families. Condensation of two benzenes sharing an edge gives naphthalene, which has a well-developed chemistry of its own. In contrast, the chemistry of condensed polyhedral borane is only being developed.[1] Among the possible condensation products of $B_{12}H_{12}^{2-}$ (1 in Figure 1), such as the edge sharing $B_{22}H_{20}^{2-}$ (2 in Figure 1), face sharing $B_{21}H_{18}^{-}$ (3 in Figure 1) and four-atom sharing $B_{20}H_{16}$ (4 in Figure 1), the latter is synthesized and characterized.[2] A major part of the development in the chemistry of polyhedral boranes came from the study of carboranes. We compare here the isomers of $C_2B_{10}H_{12}$ and $CB_{19}H_{16}^{+}$. We find that such cross comparisons between two and three dimensional structures are very useful. We also study here the structure and stability of the chloroderivatives of $B_{12}H_{12}^{2-}$ and of the condensed product $B_{20}H_{16}$ and compare them to the benzenoid systems.

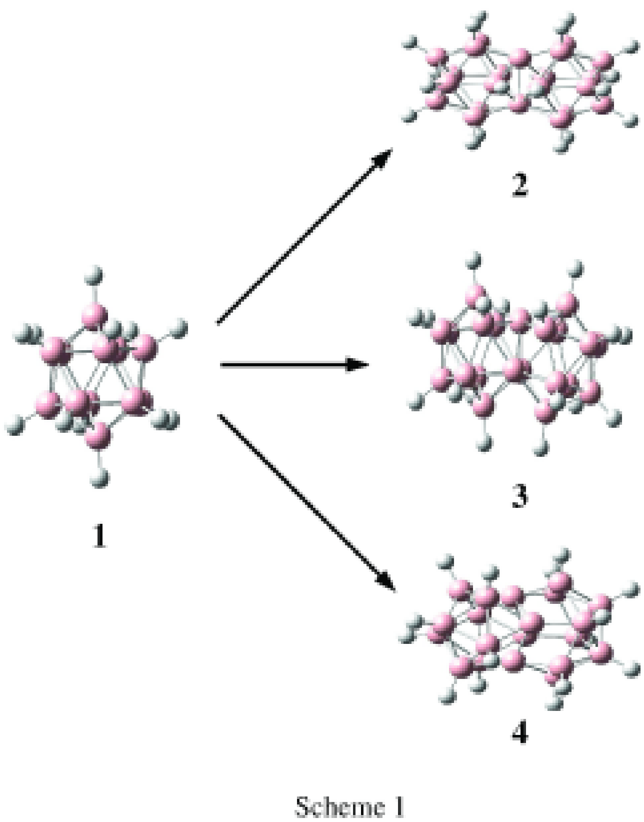

Figure 1. DFT Studies.

Methodology: All calculations used the Gaussian 03 Program packages at MHPCC.[3] Geometry optimization was carried out at B3LYP/6-31G* level of theory. All the structures are characterized by frequency analysis.

Results: We have studied the structures and relative stability of all the four mono-carborane isomers ($CB_{19}H_{16}^{+}$) of $B_{20}H_{16}$. Comparisons are made between the mono- and dichloroderivatives of two- and three-dimensional aromatic systems. We found the substituent position do not effect the extend of aromaticity significantly though there is a basic difference between two- and three-dimensional aromaticity (of having definite sigma and pi framework in the former).

References:

1) (a) J. Casanova, The Boranes, Craboranes, carborane continuum: Edition, John Wiley & Sons, 1998. (b) S. L. Shea, J. Bould, M. G. S. Londesborough, S. D. Perera, A. Franken, D. L. Ormsby, T. Jelinek, B. Stibr, J. Holub, C. A. Kilner, M. Thornton-Pett, J. D. Kennedy, *Pure Appl. Chem.*, 2003, 75, 1239. (c) E. D. Jemmis, M. M. Balakrishnarajan, P. D. Pancharatna, *J. Am. Chem. Soc.*, 2001,123,4313.
2) (a) N E. Miller, E. L. Muetterties, *J. Am. Chem. Soc*, 1963, 85, 3506. (b) L. B. Friedman, R. D. Dobrott, W. N. Lipscomb, *J. Am. Chem. Soc.*, 1963, 85, 3505. (c) N. E. Miller, C. H. Miller, E. L. Muetterties, *Inorg. Chem.,* 1964, 3, 1690. (d) R. D. Dobrott, L. B. Friedman, W. N. Lipscomb, *J. Chem. Phys.*, 1964, 40, 866.
3) M. J. Frisch *et al,* Gaussian, Inc., Pittsburgh, PA, 2003.

Author and Contact: Eluvathingal D. Jemmis
Authors: Biswarup Pathak and Anakuthil Anoop
Organization: School of Chemistry, University of Hyderabad, Hyderabad, India, 500046
Resources: IBM P3 at MHPCC
Sponsorship: Council of Scientific and Industrial Research, Department of Atomic Energy

Environment for Design of Advanced Marine Vehicles and Operations Research (ENDEAVOR) Project Tools and Architecture

Robert Dant, D.J. Fabozzi, J. Bergquist, Jeff Beck, Carl Holmberg,
Jonathan Dann, Bryan Hieda, Mary Jane Salacup, Nathan Barnes

There is a pressing need to re-evaluate the nature of naval ships and their designs. New advanced hull types can both decrease drag and detectable signature, while increasing speed and stability - providing obvious operational and budgetary pay-offs. The benefits of advanced hull designs also have a commercial spin-off in high-speed transport and ferries (see Figure 1). The ENDEAVOR project is developing a distributed architecture for assessing Advanced Marine Vehicle (AMV) performance based on operational characteristics. The Air Force Research Laboratory's Maui High Performance Computing Center (AFRL/MHPCC) is working to design, test, and integrate naval architecture tools and simulation software with oceanographic and marine vessel databases on a high performance computing system. Portions of the ENDEAVOR system are now available over the Internet for use by customers from their desktop computers.

Research Objectives: The primary objective of this research effort is to integrate a number of existing naval architecture software tools into a single, high-performance design environment. By utilizing a supercomputer server environment, the ENDEAVOR system can greatly shorten the cycle time of an iterative design and test process, enabling designers to quickly optimize complex designs. This will provide users more freedom to consider a wider range of concepts, against the time given to design a product.

Methodology: The methodology involved: 1) parameters that are manipulated by users for their designs and test conditions via a convenient Graphical User Interface (GUI), 2) a web services infrastructure (Figure 2) through which the client software submits designs to be processed and stored, 3) material and environmental databases in which to record and recall user vehicle designs and tests, and 4) physics-based codes hosted on a high performance computing system to support complex design evaluations and testing.

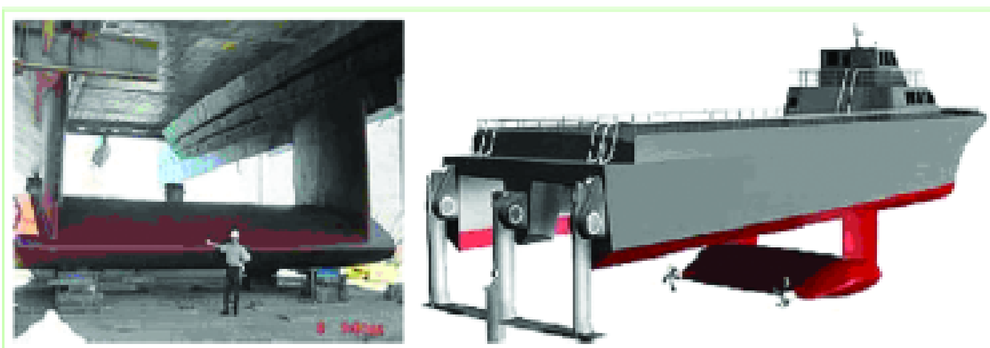

Figure 1. Sea Flyer, a Navatek Ltd designed ship (HYSWAC - Hybrid Small Waterplane Area Craft) is an example of a current Advanced Marine Vehicle (AMV).

Figure 2. Project ENDEAVOR Ocean Environment "Tools" Website.

Results: AFRL/MHPCC fielded a distributed computing infrastructure across two Linux servers and two computing clusters, for both database transactions and for complex modeling requirements. The Linux servers provided application services, including a Tomcat servlet container and SOAP (Simple Object Access Protocol) web services, as well as database management system (DBMS) services such as Oracle and MySQL.

These servers provided a suite of client-server capabilities that were designed to handle the following use-case requirements:

1) Transit Planning: answers questions regarding the best route to take, based on the particular vehicle and the anticipated near-shore and deep sea conditions.
2) Lifting Body Design: determines vehicle performance using selected design options.
3) Product Research: determines who is purchasing and/or fielding particular vehicles.
4) Harbor Operations: assists with determining improved traffic patterns.
5) Facility Construction: provides information regarding specific ocean conditions to better plan marine construction activities.
6) Amphibious Operations Support: determines near-shore conditions via powerful modeling codes and tools.
8) Environmental Modeling: investigates near-shore and deep water conditions through plots, spreadsheets, and/or Geographic Information Systems (Figure 3).

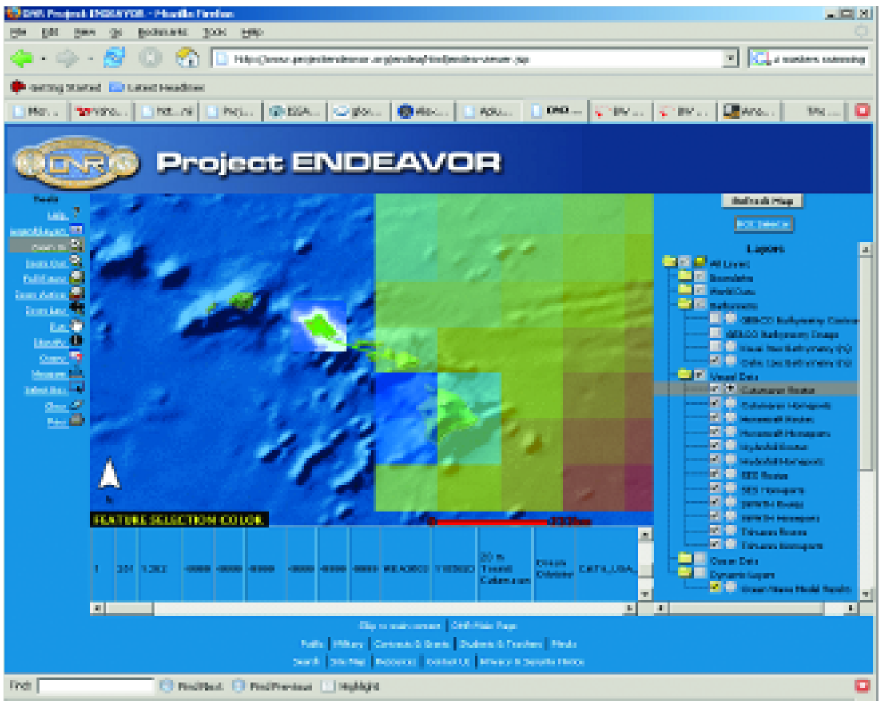

Figure 3. Project ENDEAVOR Geographic Information System (GIS).

Significance: The U.S. Navy is currently investigating a host of advanced ship design concepts that can reduce drag while increasing speed and stability. The benefits of novel hull designs also have a commercial spin-off for high-speed transports and ferries. For global transportation, cargoes can be moved from one continent to another within days instead of weeks. For regional transportation, people and cargos in offshore islands can have quick and easy access to other islands or the mainland. The impact to the global and local economies can potentially be very significant. As the ENDEAVOR system matures, designers can realistically explore new opportunities hidden in unproven design concepts.

Author and Contact: Robert Dant
Authors: D.J. Fabozzi, Jeff Beck, Carl Holmberg, Jonathan Dann, Bryan Hieda, Mary Jane Salacup, Nathan Barnes
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Author: J. Bergquist
Organization: SAIC Ship Technology Division, 134 Holiday Court, Suite 318, Annapolis, MD, 21401

Resources: A Linux application server (Dell PowerEdge 2650 Dual Processor), a linux Oracle Database server (HP ProLiant DL380 G4 Intel Xeon), a PC server, two computing clusters - a shared 128-node (256-processor) *Huinalu* Linux cluster system, and a dedicated 11-node (22-processor) ENDEAVOR Red Hat Linux Cluster system at MHPCC

Sponsorship: Office of Naval Research (ONR)

Cargo Tracking Data Warehouse System Performance Enhancements

Chad Churchwell, Aaron Steigerwald, Todd Lawson

The MHPCC team was tasked to enhance the performance of the data warehouse the Office of Naval Intelligence (ONI) uses to monitor international cargo supply chain transactions. By utilizing numerous advanced database architectural features and by leveraging the high performance computing assets at MHPCC, data loading, data retrieval, and data analysis processes were streamlined providing a faster performing, more scalable, and easier to maintain data warehouse solution. MHPCC's *Tempest* (IBM SP3/SP4) and *Huinalu* (IBM Linux Supercluster) computing resources were utilized for implementation and testing.

Research Objectives: The Office of Naval Intelligence uses a data warehouse to store and analyze cargo contents and movements in the nation's premier cargo tracking system. The MHPCC undertook the task of enhancing the original implementation with the primary objectives of improving data load and retrieval times. Using advanced technologies and techniques, MHPCC Application Engineers have developed an architecture that not only facilitates rapid data loading and retrieval, but is scalable and maintainable as well. The effort utilized MHPCC's *Tempest* and *Huinalu* computing resources for implementation and testing. The *Tempest* resource consisted of one IBM Power4 node containing 32 processors and 32 GB of RAM, and the *Huinalu* resources consisted of 32 nodes each containing dual Intel 933 processors and 1 GB of RAM.

Methodology: The ONI data warehouse effort focused on three key areas. The first involved a redesign of the data warehouse's schema and architecture, which is the foundation of the system and significantly affects performance and maintainability. The second addressed the need to improve the Extraction, Transformation, and Loading (ETL) process that executes and maintains mappings between the traditional On-Line Transaction Processing (OLTP) database source and data warehouse. The third effort focused on the loading process, which introduced an intermediate step in the form of a staging environment in order to minimize system downtime while new data is loaded into the data warehouse.

The use of MHPCC's computing resources was grouped into functional areas. The flexible *Huinalu* nodes were grouped according to need, with some used for ETL process activities and others for database access and development. The *Tempest* node was used primarily to house the numerous databases needed for development and testing.

Results: The design of a data warehouse's schema and underlying storage configuration impacts every aspect of an On-Line Analytical Processing (OLAP) system, which is a generic term for ONI's cargo tracking system, and is crucial to its success. Details such as the type of schema and configuration of the underlying storage objects must be tailored to the operational nuances of the system. The following list highlights the different technologies and architectural methods employed in implementing ONI's data warehousing solution.

Star Schema - This type of schema facilitates rapid data retrieval within the data warehouse. A detailed knowledge of how the data interacts is required for its design but supports all types of queries.

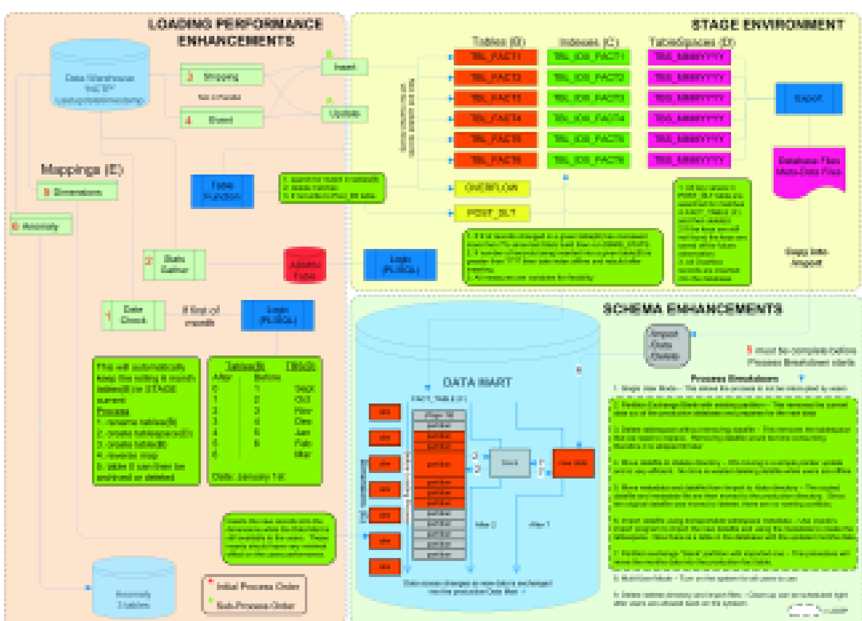

Figure 1. Architecture Overview for Cargo Tracking.

Data Partitioning & Sub Partitioning - The storage of data can be partitioned, i.e., logically and physically organized by a data field such as a date. In this instance, queries containing date criteria can significantly reduce the amount of data searched. Improved data loading and maintenance simplification is also achieved using partitioning. MHPCC engineers used a mathematically formulated date field to organize ONI's partitions.

Data Aggregation - Data sets that are commonly accessed and queried can be aggregated into objects that expedite data retrieval. ONI's data warehouse architecture takes advantage of these structures, improving query response time.

Database Management System (DBMS) Optimizer Customization - A DBMS optimizer controls how every query made against the data warehouse will be executed. ONI's data warehouse uses an optimizer that was tailored to its operational characteristics and in turn improves performance.

Automated Query Regeneration - The data warehouse was configured to take advantage of a technology where by the DBMS determines whether or not to automatically rewrite the user's query in order to provide quicker results. The DBMS is able to do this based on statistics and information it gathers about the sets of data the user queries.

Bitmap Indexing - Indexing data fields improves query performance. Bitmap indexes, as opposed to b-tree indexes, are applicable when the data sets they are applied to are of a finite nature. Their use improves query performance while reducing management overhead, and have been used extensively throughout ONI's data warehouse.

There are several important aspects to selecting and implementing a successful ETL tool. Every data warehouse requires some form of ETL tool, which pulls data out of a source database, processes and manipulates it, and inserts it into a data warehouse. The data warehouse is an aggregation of the source database and has a completely different schema that is optimized for rapid data retrieval, hence the need for data processing and manipulation. The ETL tool selected for ONI's system streamlines the mappings between source and destination (data warehouse) data relationships. It also facilitates the aggregation of selected source data fields, such as multiple cargo descriptions, into one. Other technical details, such as the reuse of data record keys, proved extremely helpful during testing and validation. One critical phase of ETL tool integration is the data-mapping phase. MHPCC staff utilized internal and external subject matter expertise to efficiently and effectively map and aggregate source data fields to their data warehouse counterparts.

MHPCC engineers utilized a combination of technologies and techniques to develop an innovative solution that provides scalability and minimizes system downtime during data loading. The initial challenge was to reduce data loading times, and consequently system downtimes, that were directly proportionate to the amount of data being loaded as well as the number of records stored in the data warehouse. The first major component of the solution involved incorporating a staging environment into the system's architecture. The staging environment includes a modified version of ONI's operational data warehouse that accepts ETL inputs while the ETL is processing data, a time consuming task. The second component utilizes a data partitioning scheme and the ability to move low level database storage objects between databases, which enables the transfer of monthly blocks of data between the staging and operational data warehouses. The users experience no interruptions while the ETL loads data into the staging environment and minimal interruption while the updated storage objects are transferred into the operational data warehouse. This solution is highly scalable. Given any type of computing hardware, loading and system down time will remain constant whether 50 K or 50 M records are loaded per day. As a result, users will never experience fluctuations in performance and downtime is minimal.

The implementation and testing of all data warehouse enhancements were made possible utilizing MHPCC's *Tempest* and *Huinalu* computing resources. *Tempest*, a 32 IBM Power4 processor, 32 GB RAM node, housed virtually all databases used during development and testing. It allowed engineers to experiment with different techniques requiring varying degrees of processing power. The *Huinalu* nodes, each containing two Intel 933 processors and 1 GB RAM, executed ETL processes and served as a development environment for our geographically separated team members. The ability to easily reconfigure different clusters of *Huinalu* nodes afforded the team the flexibility it needed during different stages of the project.

Significance: ONI's premier cargo tracking system's data warehouse has undergone significant enhancements at the MHPCC. Its schema and system architecture have been modified to improve data loading and retrieval times using a variety of cutting edge technologies and techniques. An ETL tool was selected and implemented that streamlines mappings and eases their future maintenance. The data loading solution affords improved scalability and reduces system downtime. ONI's analysts benefit from the enhancements by having consistent and fast access to crucial data, while ONI's information technology support staff benefit from improved system scalability and maintainability.

Author and Contact: Chad Churchwell
Authors: Aaron Steigerwald and Todd Lawson
Organization: Maui High Performance Computing Center, 550 Lipoa Parkway, Kihei, HI, 96753
Resources: *Tempest* (IBM SP3/SP4) and *Huinalu (*IBM Linux Supercluster*)* at MHPCC
Sponsorship: Office of Naval Intelligence (ONI)

Evaluation of Rotorcraft Downwash Effects on Human Performance and Pararescue Operations

Ken Wurtzler

Pararescue operations performed in close proximity to rescue rotorcraft create extreme conditions for the personnel on the ground. By knowing the characteristics of the flowfield underneath rescue aircraft, the impact on the rescue personnel (PJs) can be determined, and proper training can be put into effect to prepare for actual conditions. Using the latest version of Cobalt, a modeling technique was developed using blade-element theory to model the rotor movement. Combined with DES and time-accurate simulations, a realistic picture of the flowfield was obtained. This modeling approach will be used for future vehicle development as defined by operating requirements.

Research Objectives: A computational study was performed on several notional aircraft at various heights to determine radial velocities and dynamic pressures acting on rescue personnel approaching the vehicle or entering/exiting the vehicle via ahoist. A total of 22 cases were completed - [4 configurations at 4 heights, 4 cases with PJs, and 2 headwind cases]. Time-accurate data was calculated over a particular time-span and then post-processed to locate peak radial velocities and pressure data on the rescue personnel. With several of the extremes reaching twice the time-averaged value, it was paramount to gather time-accurate data in order to understand in what conditions rescue personnel operate. This was substantiated by a corresponding experimental ground test with rescue personnel.

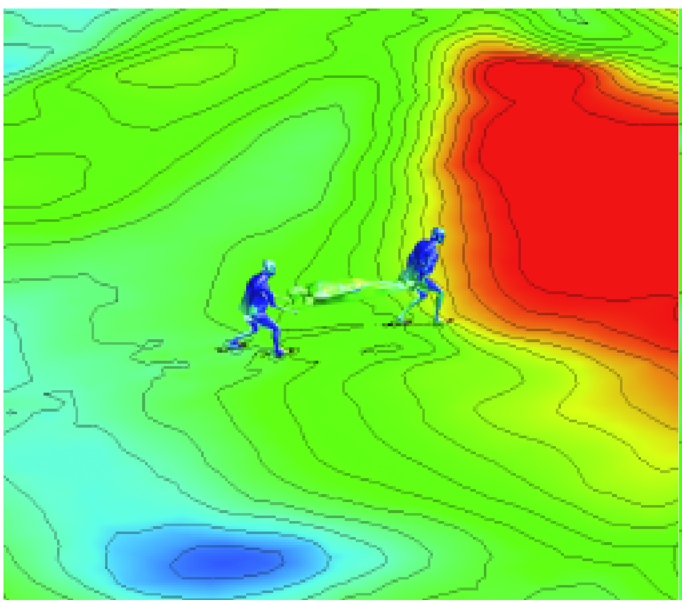

Figure 1. PJs approaching downwash under notional aircraft.

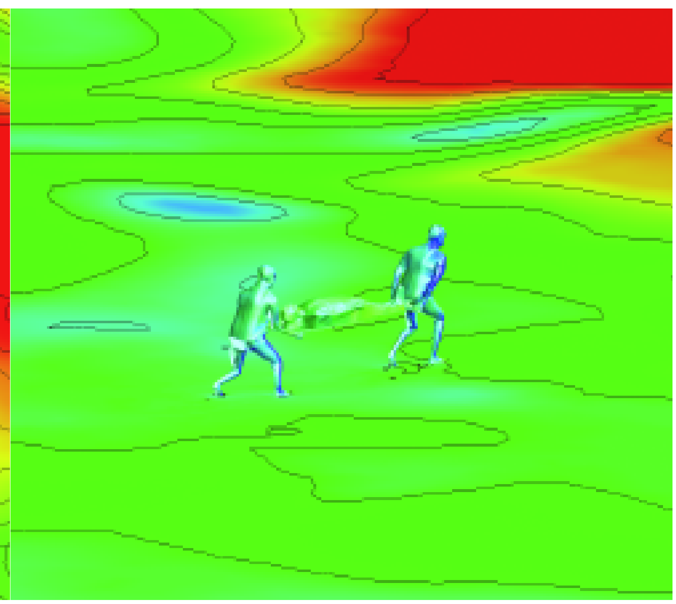

Figure 2. PJs approaching downwash under notional aircraft.

Author: Ken Wurtzler
Organization: Cobalt Solutions, LLC, 729 Miami Heights, Loveland, Ohio, 45140
Resources: 256 Processors on *Tempest* (IBM SP3/SP4) at MHPCC
Sponsorship: Funding for this program is from the Combat Search and Rescue (CSAR) SPO, SOFSG/TH

INDEX OF AUTHORS

MAUI HIGH PERFORMANCE COMPUTING CENTER

550 Lipoa Parkway, Kihei, Hawaii 96753

Phone: 808-879-5077 - Fax: 808-879-5018

Email: info@mhpcc.hpc.mil

http: //www.mhpcc.hpc.mil